

IBM WebSphere Application Server Enterprise,  
Version 5.0.2



# Applications

**Note**

Before using this information, be sure to read the general information under “Trademarks and service marks” on page xiii.

**Compilation date: July 22, 2003**

**© Copyright International Business Machines Corporation 2003. All rights reserved.**

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

# Contents

|  |      |
|--|------|
| Trademarks and service marks . . . . . | xiii |
|--|------|

|  |   |
|--|---|
| Chapter 1. Welcome to Applications . . . . . | 1 |
|--|---|

|  |    |
|--|----|
| Chapter 2. Using Web applications. . . . . | 11 |
|--|----|

|  |    |
|--|----|
| Web applications . . . . .   | 11 |
| web.xml file . . . . .   | 12 |
| Migrating Web application components . . . . .   | 13 |
| Default Application. . . . .   | 16 |
| Snoop . . . . .  | 16 |
| HelloHTML . . . . .  | 17 |
| HitCount . . . . .   | 17 |
| Servlets. . . . .  | 17 |
| Developing servlets with WebSphere Application Server extensions . . . . .                 | 18 |
| Application lifecycle listeners and events . . . . .                                       | 18 |
| Listener classes for servlet context and session changes. . . . .                          | 19 |
| Example:   |    |
| com.ibm.websphere.DBConnectionListener.java . . . . .                                      | 19 |
| Servlet filtering . . . . .  | 20 |
| Filter, FilterChain, FilterConfig classes for servlet filtering. . . . .                   | 20 |
| Example: com.ibm.websphere.LoggingFilter.java . . . . .                                    | 20 |
| Configuring page list servlet client configurations . . . . .                              | 21 |
| autoRequestEncoding and autoResponseEncoding . . . . .                                     | 25 |
| autoRequestEncoding and autoResponseEncoding encoding examples . . . . .                   | 25 |
| JavaServer Pages files . . . . .   | 26 |
| Developing JavaServer Pages files with WebSphere extensions . . . . .                      | 26 |
| Tag libraries . . . . .  | 27 |
| tsx:dbconnect tag JavaServer Pages syntax . . . . .  | 27 |
| dbquery tag JavaServer Pages syntax . . . . .  | 28 |
| dbmodify tag JavaServer Pages syntax . . . . .   | 29 |
| tsx:getProperty tag JavaServer Pages syntax and examples . . . . .                         | 30 |
| tsx:userid and tsx:passwd tag JavaServer Pages syntax . . . . .                            | 31 |
| tsx:repeat tag JavaServer Pages syntax . . . . .   | 31 |
| Example: Combining tsx:repeat and tsx:getProperty tag JavaServer Pages tags. . . . .       | 32 |
| Example: tsx:dbmodify tag syntax . . . . .   | 32 |
| Example: Using tsx:repeat tag JavaServer Pages tag to iterate over a results set . . . . . | 32 |
| JspBatchCompiler tool. . . . .   | 35 |
| Bean Scripting Framework . . . . .   | 36 |
| Developing Web applications . . . . .  | 36 |
| Disabling JavaServer Pages run-time compilation . . . . .                                  | 37 |
| Example: Converting JavaScript source to the Bean Scripting Framework . . . . .            | 38 |
| Scenario: Creating a Bean Scripting Framework application . . . . .                        | 39 |
| Scenario description . . . . .   | 39 |
| Developing the BSF application. . . . .  | 40 |

|   |    |
|---|----|
| Deploying the BSF application . . . . .   | 44 |
| Example: Bean Scripting Framework code example . . . . .                        | 45 |
| Web modules. . . . .  | 48 |
| Assembling Web Modules . . . . .  | 48 |
| Context parameters. . . . .   | 50 |
| Security constraints. . . . .   | 50 |
| Servlet mappings . . . . .  | 51 |
| Invoker attributes . . . . .  | 51 |
| Error pages . . . . .   | 51 |
| File serving . . . . .  | 51 |
| Initialization parameters . . . . .   | 51 |
| Servlet caching . . . . .   | 51 |
| Web components . . . . .  | 51 |
| Web property extensions . . . . .   | 52 |
| Web resource collections . . . . .  | 52 |
| Welcome files. . . . .  | 52 |
| Context parameter assembly settings . . . . .                                   | 52 |
| Initialization parameter assembly settings . . . . .                            | 52 |
| Filter assembly settings . . . . .  | 53 |
| JavaServer Pages attribute assembly settings . . . . .                          | 53 |
| Multipurpose Internet Mail Extensions (MIME) filter assembly settings . . . . . | 56 |
| Page list assembly settings . . . . .   | 57 |
| Security constraint assembly settings . . . . .                                 | 58 |
| Servlet mapping assembly settings . . . . .                                     | 59 |
| Tag library assembly settings . . . . .   | 59 |
| Welcome file assembly settings . . . . .  | 60 |
| Servlet caching configuration assembly settings . . . . .                       | 60 |
| Web components assembly settings . . . . .                                      | 62 |
| Web modules assembly settings . . . . .   | 68 |
| Assembly property extensions . . . . .  | 71 |
| File serving attribute assembly settings . . . . .                              | 72 |
| Invoker attribute assembly settings . . . . .                                   | 72 |
| Error page assembly settings . . . . .  | 72 |
| Web resource collections security constraint properties . . . . .               | 73 |
| Troubleshooting tips for Web application deployment . . . . .                   | 74 |
| Modifying the default Web container configuration . . . . .                     | 75 |
| Web container . . . . .   | 76 |
| Web container settings. . . . .   | 76 |
| Web module settings . . . . .   | 77 |
| Web Module Deployment settings. . . . .   | 77 |
| Web container custom property settings . . . . .                                | 78 |
| Web applications: Resources for learning . . . . .                              | 79 |

|   |    |
|---|----|
| Chapter 3. Managing HTTP sessions . . . . . | 81 |
|---|----|

|   |    |
|---|----|
| Sessions . . . . .                                      | 81 |
| Migrating HTTP sessions . . . . .                       | 82 |
| Developing session management in servlets . . . . .     | 83 |
| SessionSample.java . . . . .                            | 84 |
| Assembling so that session data can be shared . . . . . | 85 |
| Servlet API Behavior . . . . .                          | 85 |
| Session security support . . . . .                      | 85 |
| Security integration rules for HTTP sessions . . . . .  | 86 |

|   |     |
|---|-----|
| Programmatic details and scenarios . . . . .  | 86  |
| Session management support . . . . .  | 87  |
| Configuring session management by level . . . . .   | 88  |
| Session tracking options . . . . .  | 88  |
| Session tracking with cookies . . . . .   | 89  |
| Session tracking with URL rewriting . . . . .   | 89  |
| Session tracking with SSL information . . . . .   | 90  |
| Configuring session tracking . . . . .  | 90  |
| Serializing access to session data . . . . .  | 91  |
| Session Management settings . . . . .   | 91  |
| Cookie settings . . . . .   | 94  |
| Configuring session tracking for Wireless<br>Application Protocol (WAP) devices . . . . . | 95  |
| Session management custom properties . . . . .  | 96  |
| Distributed sessions . . . . .  | 97  |
| Session recovery support . . . . .  | 97  |
| Distributed Environment settings . . . . .  | 97  |
| Configuring for database session persistence . . . . .                                    | 98  |
| Switching to a multirow schema . . . . .  | 98  |
| Configuring tablespace and page sizes for DB2<br>session databases . . . . .              | 98  |
| Database settings . . . . .   | 99  |
| Multirow schema considerations . . . . .  | 100 |
| Clustered session support . . . . .   | 101 |
| Tuning session management . . . . .   | 102 |
| Configuring scheduled invalidation . . . . .  | 102 |
| Configuring write contents . . . . .  | 103 |
| Configuring write frequency . . . . .   | 104 |
| Base in-memory session pool size . . . . .  | 104 |
| Controlling write operations . . . . .  | 105 |
| Tuning parameter settings . . . . .   | 105 |
| Best practices for using HTTP Sessions . . . . .  | 107 |
| Managing HTTP sessions: Resources for learning: . . . . .                                 | 110 |

## Chapter 4. Using enterprise beans in applications . . . . . 113

|   |     |
|---|-----|
| Enterprise beans . . . . .  | 113 |
| Developing enterprise beans . . . . .                                       | 114 |
| Migrating enterprise bean code to the supported<br>specification . . . . .  | 115 |
| WebSphere extensions to the Enterprise<br>JavaBeans specification . . . . . | 118 |
| Best practices for developing enterprise beans . . . . .                    | 119 |
| Using access intent policies . . . . .                                      | 124 |
| Access intent policies . . . . .  | 124 |
| Access intent service . . . . .   | 127 |
| Access intent design considerations . . . . .                               | 128 |
| Applying access intent policies to methods . . . . .                        | 128 |
| Using the AccessIntent API . . . . .  | 129 |
| Access intent exceptions . . . . .  | 131 |
| Access intent assembly settings . . . . .                                   | 132 |
| Access intent best practices . . . . .                                      | 134 |
| Frequently asked questions: Access intent . . . . .                         | 134 |
| EJB modules . . . . .   | 136 |
| Assembling EJB modules . . . . .  | 136 |
| CMP field assembly settings . . . . .                                       | 138 |
| Container transactions . . . . .  | 138 |
| Container transaction assembly settings . . . . .                           | 138 |
| EJB module assembly settings . . . . .                                      | 140 |
| Entity bean assembly settings . . . . .                                     | 141 |
| Message-driven bean assembly settings . . . . .                             | 151 |

|  |     |
|--|-----|
| Method extensions . . . . .                        | 155 |
| Method extension assembly settings . . . . .       | 156 |
| Method permissions . . . . .                       | 158 |
| Method permission assembly settings . . . . .      | 159 |
| Query assembly settings . . . . .                  | 159 |
| References . . . . .                               | 160 |
| EJB reference assembly settings . . . . .          | 160 |
| EJB local-reference assembly settings . . . . .    | 161 |
| EJB relation assembly settings . . . . .           | 162 |
| Exclude list assembly settings . . . . .           | 162 |
| Security role assembly settings . . . . .          | 163 |
| Session bean assembly properties . . . . .         | 164 |
| EJB containers . . . . .                           | 169 |
| Managing EJB containers . . . . .                  | 170 |
| EJB container settings . . . . .                   | 170 |
| EJB container system properties . . . . .          | 171 |
| EJB cache settings . . . . .                       | 172 |
| Container interoperability . . . . .               | 173 |
| Deploying EJB modules . . . . .                    | 178 |
| EJB module collection . . . . .                    | 178 |
| EJB module settings . . . . .                      | 179 |
| Enterprise beans: Resources for learning . . . . . | 179 |
| EJB method Invocation Queuing . . . . .            | 181 |

## Chapter 5. Using extended messaging in applications . . . . . 185

|   |     |
|---|-----|
| Extended messaging - overview . . . . .                                     | 185 |
| Extended messaging - receiving messages . . . . .                           | 186 |
| Extended messaging - sending messages . . . . .                             | 188 |
| Extended messaging - data mapping . . . . .                                 | 189 |
| Extended messaging - handling late responses . . . . .                      | 189 |
| Extended messaging - transactional support . . . . .                        | 190 |
| Extended messaging - exception handling . . . . .                           | 192 |
| Extended messaging - application usage scenarios . . . . .                  | 193 |
| Extended messaging - components . . . . .                                   | 194 |
| Designing an enterprise application to use<br>extended messaging . . . . .  | 196 |
| Developing an enterprise application to use<br>extended messaging . . . . . | 197 |
| Deploying an enterprise application to use<br>extended messaging . . . . .  | 198 |
| Configuring deployment attributes for a receiver<br>bean . . . . .          | 199 |
| Configuring deployment attributes for a sender<br>bean . . . . .            | 201 |
| Configuring extended messaging service resources . . . . .                  | 202 |
| Adding a new input port . . . . .   | 202 |
| Adding a new output port . . . . .  | 203 |
| Configuring an input port . . . . .   | 203 |
| Configuring an output port . . . . .  | 204 |
| Extended messaging service settings . . . . .                               | 204 |
| Extended messaging provider settings . . . . .                              | 206 |
| Troubleshooting extended messaging . . . . .                                | 214 |
| Extended Messaging: Resources for learning . . . . .                        | 215 |

## Chapter 6. Using message-driven beans in applications . . . . . 217

|  |     |
|--|-----|
| Message-driven beans - an overview . . . . .         | 217 |
| Message-driven beans - components . . . . .          | 218 |
| Message-driven beans - transaction support . . . . . | 220 |

|  |     |
|--|-----|
| Designing an enterprise application to use message-driven beans . . . . .  | 220 |
| Developing an enterprise application to use message-driven beans . . . . . | 222 |
| Migrating a JMS listener application to use message-driven beans . . . . . | 224 |
| Deploying an enterprise application to use message-driven beans . . . . .  | 225 |
| Configuring deployment attributes for a message-driven bean . . . . .      | 226 |
| Configuring message listener resources for message-driven beans . . . . .  | 228 |
| Configuring the message listener service . . . . .                         | 228 |
| Adding a new listener port . . . . .                                       | 233 |
| Configuring a listener port . . . . .                                      | 234 |
| Deleting a listener port . . . . .   | 234 |
| Configuring security for message-driven beans . . . . .                    | 235 |
| Administering listener ports . . . . .                                     | 235 |
| Important files for message-driven beans and extended messaging . . . . .  | 237 |
| Troubleshooting message-driven beans . . . . .                             | 237 |
| Message-driven beans samples . . . . .                                     | 238 |

## Chapter 7. Using application clients **241**

|  |     |
|--|-----|
| Application clients . . . . .  | 241 |
| Application client functions . . . . .   | 243 |
| ActiveX application clients . . . . .  | 245 |
| Applet clients . . . . .   | 246 |
| J2EE application clients . . . . .   | 247 |
| Pluggable application clients . . . . .  | 248 |
| Thin application clients . . . . .   | 249 |
| Example: Migrating application clients . . . . .   | 251 |
| Migration tips for application clients . . . . .   | 252 |
| Installing application clients . . . . .   | 252 |
| Installing application clients on Version 9 of the Solaris Operating Environment . . . . . | 254 |
| Developing ActiveX application client code . . . . .                                       | 256 |
| Starting an ActiveX application . . . . .  | 256 |
| JClassProxy and JObjectProxy classes . . . . .   | 259 |
| Java virtual machine initialization tips . . . . .   | 263 |
| Example: Developing ActiveX to enterprise bean bridge, using Java proxy objects . . . . .  | 263 |
| Example: Calling Java methods in the ActiveX to enterprise bean bridge . . . . .           | 264 |
| Java field programming tips . . . . .  | 266 |
| ActiveX to Java primitive data type conversion values . . . . .                            | 266 |
| Array tips for ActiveX application clients . . . . .                                       | 268 |
| Error handling codes for ActiveX application clients . . . . .                             | 268 |
| Threading tips . . . . .   | 269 |
| Example: Viewing System.out message . . . . .  | 270 |
| Example: Enabling logging and tracing for application clients . . . . .                    | 271 |
| ActiveX client programming best practices . . . . .  | 272 |
| Developing applet client code . . . . .  | 275 |
| Accessing secure resources using the TCP/IP protocol for applet clients . . . . .          | 276 |
| Applet client tag requirements . . . . .   | 277 |
| Applet client code requirements . . . . .  | 277 |
| Developing J2EE application client code . . . . .  | 278 |

|  |     |
|--|-----|
| J2EE application client class loading . . . . .  | 281 |
| Developing pluggable application client code. . . . .  | 283 |
| Developing thin application client code. . . . .   | 284 |
| Developing thin application client code on a client machine . . . . .  | 284 |
| Developing thin application client code on a server machine . . . . .  | 285 |
| Assembling Application Client Modules . . . . .  | 286 |
| Application client modules . . . . .   | 287 |
| Application client assembly settings . . . . .   | 288 |
| Environment entries assembly properties . . . . .  | 289 |
| Deploying application clients on workstation platforms. . . . .  | 290 |
| Starting the Application Client Resource Configuration Tool and opening an EAR file . . . . .  | 291 |
| Data sources for application clients . . . . .   | 291 |
| Configuring new data source providers (JDBC providers) for application clients . . . . .   | 291 |
| Configuring new data sources for application clients . . . . .   | 294 |
| Mail providers and mail sessions for the Application Client Assembly Tool . . . . .  | 295 |
| Configuring mail providers and sessions for application clients . . . . .  | 295 |
| Configuring new mail sessions for application clients . . . . .  | 297 |
| URLs for application clients . . . . .   | 298 |
| URL providers for the Application Client Resource Configuration Tool . . . . .   | 298 |
| Configuring new URL providers for application clients . . . . .  | 298 |
| Configuring new URLs with the Application Client Resource Configuration Tool . . . . .   | 301 |
| WebSphere asynchronous messaging using the Java Message Service API for the Application Client Resource Configuration Tool . . . . . | 301 |
| Configuring Java messaging client resources . . . . .  | 302 |
| Configuring new connection factories for application clients . . . . .   | 332 |
| Configuring new Java Message Service destinations for application clients . . . . .  | 332 |
| Example: Configuring MQ Queue and Topic connection factories and destination factories for application clients . . . . .             | 333 |
| Example: Configuring WAS Queue and Topic connection factories and destination factories for application clients . . . . .            | 334 |
| Configuring new resource environment providers for application clients . . . . .   | 336 |
| Configuring new resource environment entries for application clients. . . . .  | 336 |
| Managing application clients . . . . .   | 337 |
| Updating data source and data source provider configurations with the Application Client Resource Configuration Tool . . . . .       | 338 |
| Updating URLs and URL provider configurations for application clients . . . . .  | 338 |
| Updating mail session configurations for application clients . . . . .   | 339 |

|  |            |
|--|------------|
| Updating Jave Message Service provider, connection factories, and destination configurations for application clients . . . . .           | 339        |
| Updating MQ Java Message Service provider, MQ connection factories, and MQ destination configurations for application clients . . . . .  | 340        |
| Updating Resource Environment Entry and Resource Environment Provider configurations for application clients. . . . .                    | 340        |
| Removing application client resources . . . . .  | 341        |
| Running application clients. . . . .   | 342        |
| launchClient tool . . . . .  | 343        |
| Example: Using a Java 2 security manager with a J2EE application client. . . . .   | 347        |
| Example: Enabling Java 2 security prior to J2EE application client runtime initialization. . . . .                                       | 347        |
| Application client troubleshooting tips . . . . .  | 348        |
| <b>Chapter 8. Using Web services . . . . .</b>   | <b>353</b> |
| Web services . . . . .   | 354        |
| Planning to use Web services . . . . .   | 355        |
| Setting up a Web services development environment. . . . .   | 356        |
| Migrating Apache SOAP Web services to Web services for J2EE . . . . .  | 357        |
| Developing Web services . . . . .  | 360        |
| Developing a Web service using a Java bean . . . . .   | 361        |
| Developing a Web service using a stateless session enterprise bean . . . . .   | 381        |
| Developing a new Web service with an existing Web Services Description Language file using a Java bean. . . . .                          | 381        |
| Developing a new Web service with an existing Web Services Description Language file using a stateless session enterprise bean . . . . . | 388        |
| Web services development artifacts . . . . .   | 389        |
| Mapping between Java, Web Services Description Language and XML . . . . .  | 390        |
| Developing a Web services client . . . . .   | 414        |
| Assembling a Web services-enabled client JAR and EAR file . . . . .  | 415        |
| Testing Web services-enabled clients. . . . .  | 417        |
| Assembling Web services applications . . . . .   | 417        |
| Assembling a Web services-enabled EJB JAR file   | 418        |
| Assembling Web services-enabled WAR file . . . . .   | 421        |
| Assembling a Web services-enabled EAR file   | 423        |
| Web services assembly properties. . . . .  | 424        |
| Enabling the EAR file . . . . .  | 428        |
| Deploying Web services . . . . .   | 432        |
| wsdeploy command . . . . .   | 433        |
| Using Java Messaging Service to transport Web services requests . . . . .  | 435        |
| Java Messaging Service endpoint URL syntax   | 437        |
| Securing Web services . . . . .  | 438        |
| Configuring client-side transport level security   | 439        |
| Configuring HTTP basic authentication. . . . .   | 441        |
| Web Services: Default bindings for the Web Services Security collection . . . . .  | 442        |
| Web Services: Server Security Bindings collection . . . . .  | 465        |
| Web Services: Client Security Bindings collection  | 471        |

|  |     |
|--|-----|
| Tuning Web services applications. . . . .  | 475 |
| Troubleshooting Web services . . . . .   | 475 |
| Tracing Web services messages . . . . .  | 476 |
| Frequently asked questions about Web services for J2EE . . . . .                         | 476 |
| Web services: Resources for learning. . . . .  | 479 |
| Web services implementation scope . . . . .  | 482 |
| Port . . . . .   | 483 |
| Service . . . . .  | 483 |
| URI . . . . .  | 483 |
| Scope . . . . .  | 483 |
| Web services client bindings . . . . .   | 483 |
| Web Service . . . . .  | 483 |
| URI . . . . .  | 483 |
| WSDL Filename . . . . .  | 483 |
| Default Port Mappings . . . . .  | 483 |
| Default Port Mapping Definitions collection . . . . .                                    | 484 |
| Port Type Local Name . . . . .   | 484 |
| Port Type Namespace . . . . .  | 484 |
| Default Port Local Name . . . . .  | 484 |
| Default Port Namespace. . . . .  | 484 |
| Default Port Type Mapping Properties settings . . . . .                                  | 484 |
| Port Type Local Name . . . . .   | 484 |
| Port Type Namespace . . . . .  | 484 |
| Default Port Local Name . . . . .  | 485 |
| Default Port Namespace. . . . .  | 485 |
| Publish WSDL files settings . . . . .  | 485 |
| HTTP . . . . .   | 485 |
| Select HTTP URL prefix . . . . .   | 485 |
| Custom HTTP URL prefix . . . . .   | 485 |
| JMS . . . . .  | 485 |
| Using Apache SOAP Web services in Version 5.0 and 5.0.1 . . . . .                        | 486 |
| Developing an Apache SOAP client . . . . .   | 486 |
| Accessing enterprise beans with Apache SOAP  | 488 |
| Assembling Apache SOAP Web services . . . . .  | 488 |
| Apache SOAP deployment descriptor . . . . .  | 489 |
| Enabling Apache SOAP Web services in an enterprise application . . . . .                 | 492 |
| Deploying Apache SOAP Web services applications . . . . .                                | 494 |
| Administering deployed Apache SOAP Web services (XML-SOAP administrative tool) . . . . . | 494 |
| Securing Apache SOAP Web services . . . . .  | 495 |
| UDDI4J specifications . . . . .  | 499 |
| Web services: Resources for learning. . . . .  | 499 |

|  |            |
|--|------------|
| <b>Chapter 9. Web Services Invocation Framework (WSIF): Enabling Web services. . . . .</b> | <b>503</b> |
| Goals of WSIF . . . . .  | 503        |
| WSIF - Web services are not just SOAP services   | 504        |
| WSIF - tying client code to a particular protocol implementation is restricting . . . . .  | 504        |
| WSIF - incorporating new bindings into client code is hard . . . . .                       | 504        |
| WSIF - multiple bindings can be used in flexible ways . . . . .                            | 504        |
| WSIF - a freer Web services environment enables intermediaries . . . . .                   | 505        |

|  |     |
|--|-----|
| An overview of WSIF . . . . .  | 505 |
| WSIF architecture . . . . .  | 505 |
| Using WSIF with Web services that offer multiple bindings . . . . .    | 506 |
| WSIF and WSDL . . . . .  | 506 |
| WSIF usage scenarios. . . . .  | 507 |
| Dynamic invocation . . . . .   | 508 |
| Using WSIF to invoke Web services . . . . .                            | 508 |
| Using the WSIF providers . . . . .                                     | 509 |
| Developing a WSIF service . . . . .                                    | 521 |
| Using complex types . . . . .  | 530 |
| Using JNDI . . . . .   | 531 |
| Passing SOAP messages with attachments using WSIF . . . . .            | 533 |
| Interacting with the WebSphere J2EE container . . . . .                | 536 |
| Running WSIF as a client . . . . .                                     | 536 |
| WSIF system management and administration . . . . .                    | 536 |
| Maintaining the WSIF properties file . . . . .                         | 536 |
| Enabling security for WSIF . . . . .                                   | 537 |
| WSIF troubleshooting tips . . . . .                                    | 538 |
| WSIF API . . . . .   | 542 |
| WSIF API reference: Creating a message for sending to a port . . . . . | 543 |
| WSIF API reference: Finding a port factory or service. . . . .         | 544 |
| WSIF API reference: Using ports . . . . .                              | 545 |
| WSIF: Resources for learning . . . . .                                 | 549 |

## **Chapter 10. IBM WebSphere UDDI Registry. . . . . 551**

|   |     |
|---|-----|
| UDDI Registry terminology . . . . .   | 551 |
| UDDI Registry definitions . . . . .   | 552 |
| An overview of IBM UDDI Registries . . . . .  | 553 |
| Migrating from the IBM WebSphere UDDI Registry on WebSphere Application Server 4.0 . . . . .                      | 554 |
| Installing and setting up a UDDI Registry . . . . .   | 555 |
| Installing the UDDI Registry into a deployment manager cell. . . . .  | 557 |
| Setting up the UDDI Registry to use Cloudscape within a deployment manager cell . . . . .                         | 560 |
| Setting up the UDDI Registry to use DB2 within a deployment manager cell. . . . .                                 | 562 |
| Installing the UDDI Registry into a single appserver . . . . .  | 565 |
| Setting up the UDDI Registry to use Cloudscape in a single appserver . . . . .                                    | 566 |
| Setting up the UDDI Registry to use DB2 in a single appserver . . . . .   | 568 |
| Reinstalling the UDDI Registry application . . . . .  | 572 |
| Applying Service to the UDDI Registry in a Network Deployment and single Application Server environment . . . . . | 572 |
| Removing the UDDI Registry application from a deployment manager cell . . . . .                                   | 573 |
| Removing the UDDI Registry application from a single appserver . . . . .  | 573 |
| Configuring the UDDI Registry . . . . .   | 574 |
| Configuring global UDDI properties. . . . .   | 574 |
| Modifying the database userid and password . . . . .  | 576 |
| Configuring security properties . . . . .   | 577 |

|   |     |
|---|-----|
| Configuring the UDDI User Console (GUI) for multiple language encoding support. . . . . | 577 |
| Customizing the UDDI User Console (GUI) . . . . .                                       | 577 |
| Configuring SOAP interface properties . . . . .   | 578 |
| Configuring SOAP properties with the AAT . . . . .                                      | 578 |
| Configuring SOAP properties in an already-deployed application . . . . .                | 579 |
| Configuring WebSphere to use HTTPS and SSL . . . . .                                    | 579 |
| Administering the UDDI Registry . . . . .   | 579 |
| Running the UDDI Registry . . . . .   | 579 |
| Backing up and restoring the UDDI Registry database . . . . .                           | 579 |
| UDDI user console . . . . .   | 580 |
| Displaying the user console . . . . .   | 584 |
| Custom Taxonomy Support in the UDDI Registry . . . . .                                  | 585 |
| SOAP Application Programming Interface for the UDDI Registry . . . . .                  | 593 |
| Programming the UDDI SOAP API . . . . .   | 594 |
| SOAP API error handling tips in the UDDI Registry . . . . .                             | 594 |
| UDDI Registry Application Programming Interface . . . . .                               | 594 |
| Inquiry API for the UDDI Registry . . . . .   | 594 |
| Browse pattern for the UDDI Registry . . . . .  | 595 |
| Drilldown pattern for the UDDI Registry . . . . .                                       | 595 |
| Invocation pattern for the UDDI Registry . . . . .                                      | 595 |
| Inquiry API functions in the UDDI Registry . . . . .                                    | 596 |
| Accessible query values in the UDDI Registry . . . . .                                  | 597 |
| Publish API for the UDDI Registry . . . . .   | 598 |
| UDDI EJB Interface for the UDDI Registry . . . . .                                      | 599 |
| Datatypes package in the UDDI Registry . . . . .  | 605 |
| EJB interface methods in the UDDI Registry . . . . .                                    | 607 |
| UDDI troubleshooting tips . . . . .   | 608 |
| Turning on UDDI trace . . . . .   | 610 |
| Messages . . . . .  | 611 |
| UDAI (Web Services UDDI) messages . . . . .   | 612 |
| UDCF (Web Services UDDI) messages . . . . .   | 612 |
| UDDA (Web Services UDDI) messages . . . . .   | 613 |
| UDDM (Web Services UDDI) messages. . . . .  | 613 |
| UDEJ (Web Services UDDI) messages . . . . .   | 613 |
| UDEX (Web Services UDDI) messages . . . . .   | 613 |
| UDIN (Web Services UDDI) messages . . . . .   | 613 |
| UDLC (Web Services UDDI) messages . . . . .   | 639 |
| UDPR (Web Services UDDI) messages . . . . .   | 639 |
| UDRS (Web Services UDDI) messages . . . . .   | 639 |
| UDSC (Web Services UDDI) messages . . . . .   | 639 |
| UDSP (Web Services UDDI) messages . . . . .   | 639 |
| UDUC (Web Services UDDI) messages . . . . .   | 641 |
| UDUU (Web Services UDDI) messages . . . . .   | 643 |
| Running the UDDI Samples . . . . .  | 643 |
| Installation Verification Program (IVP) . . . . .                                       | 643 |
| Reporting problems with the IBM WebSphere UDDI Registry . . . . .                       | 645 |
| Feedback . . . . .  | 646 |

## **Chapter 11. Web services gateway: Enabling Web services . . . . . 647**

|   |     |
|---|-----|
| Web services gateway - Frequently Asked Questions . . . . . | 647 |
| What are Web services? . . . . .                            | 648 |
| What is the IBM Web services gateway? . . . . .             | 648 |
| How does the Web services gateway work? . . . . .           | 648 |

|   |            |
|---|------------|
| What problems are solved by the Web services gateway? . . . . .                             | 648        |
| Who should use the Web services gateway? . . . . .  | 649        |
| What is the difference between the Apache SOAP channel and the SOAP/HTTP channel? . . . . . | 649        |
| Web services gateway - What is new in this release . . . . .                                | 649        |
| Web services gateway - Completing the installation . . . . .                                | 650        |
| Web services gateway - prerequisites and constraints . . . . .                              | 651        |
| Establishing requirements for using a database with the gateway . . . . .                   | 652        |
| Installing the gateway into a deployment manager cell. . . . .                              | 653        |
| Installing the gateway into a stand-alone application server . . . . .                      | 656        |
| Testing the Web services gateway installation . . . . .                                     | 659        |
| Backing up and restoring a gateway configuration . . . . .                                  | 660        |
| Backing up and restoring UDDI publication links . . . . .                                   | 662        |
| Creating and updating a gateway cluster . . . . .   | 663        |
| Administering the Web services gateway . . . . .  | 664        |
| Setting the namespace URI and WSDL URI for the Web services gateway . . . . .               | 666        |
| Working with channels . . . . .   | 667        |
| Working with filters . . . . .  | 672        |
| Working with UDDI references . . . . .  | 674        |
| Working with Web services. . . . .  | 679        |
| Running the Web services gateway samples . . . . .  | 690        |
| Passing SOAP messages with attachments through the Web services gateway . . . . .           | 690        |
| SOAP messages with attachments - a definition . . . . .                                     | 691        |
| Writing the WSDL extensions for SOAP messages with attachments. . . . .                     | 691        |
| Developing Web services gateway extensions. . . . .   | 692        |
| Writing a filter for the Web services gateway . . . . .                                     | 693        |
| Using a filter to select a target service and port . . . . .                                | 701        |
| Capturing Web service invocation information from the Web services gateway . . . . .        | 703        |
| Handling exceptions for the Web services gateway . . . . .                                  | 704        |
| Administering security for the Web services gateway . . . . .                               | 705        |
| Enabling Web Services Security (WS-Security) for the gateway. . . . .                       | 705        |
| Enabling basic authentication and authorization for the gateway. . . . .                    | 715        |
| Invoking Web services over HTTPS . . . . .  | 720        |
| Enabling proxy authentication for the gateway . . . . .                                     | 720        |
| Web services gateway troubleshooting tips . . . . .   | 722        |
| Web services gateway messages . . . . .   | 726        |
| Web services gateway: Resources for learning . . . . .                                      | 738        |
| <b>Chapter 12. Class loading . . . . .</b>  | <b>739</b> |
| Class loaders . . . . .   | 740        |
| Class loader collection . . . . .   | 743        |
| Classloader ID . . . . .  | 744        |
| Classloader Mode . . . . .  | 744        |
| Class loader settings . . . . .   | 744        |
| Migrating the class-loader Module Visibility Mode setting. . . . .                          | 744        |
| Class loading: Resources for learning . . . . .   | 745        |

|   |            |
|---|------------|
| <b>Chapter 13. Using EJB query . . . . .</b>                              | <b>747</b> |
| EJB query language . . . . .  | 747        |
| Example: EJB queries. . . . .   | 748        |
| FROM clause . . . . .   | 750        |
| Inheritance in EJB query. . . . .   | 751        |
| Path expressions . . . . .  | 752        |
| WHERE clause . . . . .  | 753        |
| Scalar functions . . . . .  | 760        |
| Aggregation functions . . . . .   | 763        |
| SELECT clause . . . . .   | 765        |
| ORDER BY clause. . . . .  | 765        |
| Subqueries . . . . .  | 766        |
| EJB query restrictions. . . . .   | 766        |
| EJB Query: Reserved words . . . . .                                       | 767        |
| EJB query: BNF syntax . . . . .   | 768        |
| Comparison of EJB 2.0 specification and WebSphere query language. . . . . | 770        |
| Using the dynamic query service. . . . .                                  | 770        |
| Example: Dynamic query remote client. . . . .                             | 772        |
| Example: Dynamic query from local client. . . . .                         | 773        |
| Dynamic query service performance considerations . . . . .                | 774        |

|   |            |
|---|------------|
| <b>Chapter 14. Using the internationalization service . . . . .</b>                   | <b>777</b> |
| Internationalization . . . . .  | 778        |
| Internationalization service: Overview . . . . .                                      | 779        |
| The internationalization service solution . . . . .                                   | 779        |
| Internationalization challenges in distributed applications . . . . .                 | 780        |
| Migrating internationalized applications . . . . .                                    | 781        |
| Assembling internationalized applications. . . . .                                    | 782        |
| Setting the internationalization type for servlets . . . . .                          | 782        |
| Configuring container internationalization for servlets . . . . .                     | 783        |
| Setting internationalization type for enterprise beans . . . . .                      | 786        |
| Configuring container internationalization for enterprise beans . . . . .             | 786        |
| Using the internationalization context API. . . . .                                   | 789        |
| Gaining access to the internationalization context API . . . . .                      | 790        |
| Accessing caller locales and time zone . . . . .                                      | 791        |
| Accessing invocation locales and time zone . . . . .                                  | 792        |
| Example: Internationalization context in an EJB client program . . . . .              | 794        |
| Example: Internationalization context in an EJB servlet. . . . .                      | 795        |
| Example: Internationalization context in a session bean. . . . .                      | 797        |
| Internationalization context API: Programming reference . . . . .                     | 799        |
| Managing the internationalization service . . . . .                                   | 810        |
| Enabling the internationalization service for servlets and enterprise beans . . . . . | 811        |
| Enabling the internationalization service for EJB clients . . . . .                   | 812        |
| Troubleshooting the internationalization service . . . . .                            | 813        |
| Internationalization service errors . . . . .   | 813        |
| Internationalization service exceptions . . . . .                                     | 816        |



|  |            |
|--|------------|
| Internationalization: Resources for learning . . .       | 816        |
| <b>Chapter 15. Application profiling . . .</b>           | <b>819</b> |
| Application profiling: Overview . . .                    | 819        |
| Tasks . . .  | 820        |
| Application profiles . . .                               | 821        |
| Application profiling performance considerations . . .   | 822        |
| Assembling applications for application profiling        | 824        |
| Applying access intent policies to entity beans          | 825        |
| Creating a custom access intent policy . . .             | 829        |
| Configuring a component task policy . . .                | 835        |
| Configuring a container task policy . . .                | 836        |
| Creating an application profile . . .                    | 838        |
| Configuring tasks on application profiles . . .          | 840        |
| Dynamic query assembly settings . . .                    | 840        |
| Managing application profiles . . .                      | 841        |
| Application profiling exceptions . . .                   | 842        |
| Application profiling service settings . . .             | 842        |
| Application profile collection . . .                     | 842        |
| Using the TaskNameManager interface . . .                | 844        |
| TaskNameManager interface . . .                          | 845        |
| <br>   |            |
| <b>Chapter 16. Using Business Rule Beans . . .</b>       | <b>847</b> |
| Advantages of externalizing business rules . . .         | 847        |
| Overview of Business Rule Beans. . .                     | 848        |
| Externalized business rules . . .                        | 849        |
| Types of business rules . . .                            | 850        |
| Rule folders . . .                                       | 851        |
| Rule attributes . . .                                    | 851        |
| Rule states . . .  | 853        |
| Rule results . . .                                       | 854        |
| Dependent rules . . .                                    | 854        |
| BRBeans run-time environment . . .                       | 855        |
| BRBeans run-time behavior. . .                           | 855        |
| BRBeans run-time exception handling . . .                | 856        |
| Rule implementors . . .                                  | 857        |
| Trigger point framework . . .                            | 858        |
| Trigger points . . .                                     | 859        |
| As Of Date . . .   | 863        |
| Predefined strategy objects . . .                        | 863        |
| Customized strategy objects . . .                        | 865        |
| Customized rule implementors . . .                       | 867        |
| Rule management command . . .                            | 868        |
| Rule importer command. . .                               | 869        |
| Rule exporter command. . .                               | 870        |
| BRBeans properties file . . .                            | 871        |
| Database considerations for BRBeans . . .                | 872        |
| Rule Management Application. . .                         | 874        |
| Rule management APIs . . .                               | 875        |
| BRBeans performance enhancements . . .                   | 876        |
| Developing BRBeans . . .                                 | 879        |
| Determining where to place a trigger point . . .         | 880        |
| Placing a trigger point in the application code          | 881        |
| Administering strategy objects to control triggers . . . | 882        |
| Implementing business rules . . .                        | 883        |
| Assembling applications for use with BRBeans . . .       | 884        |
| Managing rules. . .                                      | 886        |

|   |     |
|---|-----|
| Starting the Rule Management Application . . .    | 887 |
| Copying or moving rules or rule folders . . .     | 887 |
| Working with Quick Copy . . .                     | 888 |
| Finding a rule . . .                              | 888 |
| Deleting rules . . .                              | 889 |
| Deleting rule folders . . .                       | 889 |
| Changing the properties of a rule. . .            | 889 |
| Importing a rule . . .                            | 890 |
| Exporting a rule . . .                            | 890 |
| Renaming rules. . .                               | 891 |
| Renaming rule folders . . .                       | 891 |
| Specifying columns . . .                          | 891 |
| Changing the date and time format . . .           | 892 |
| Rule Browser . . .                                | 892 |
| File menu . . .                                   | 893 |
| Edit menu . . .                                   | 909 |
| View menu . . .                                   | 912 |
| Find Rules window . . .                           | 913 |
| Business rule beans: Resources for learning . . . | 918 |

## **Chapter 17. Using asynchronous beans. . . . . 921**

|  |     |
|--|-----|
| Asynchronous beans . . .                                 | 921 |
| Example: Asynchronous bean connection management . . .   | 923 |
| Configuring work managers . . .                          | 924 |
| Work managers. . .                                       | 925 |
| Work manager collection . . .                            | 927 |
| Work manager service settings . . .                      | 929 |
| Assembling applications that use work managers           | 929 |
| Developing work objects to run code in parallel          | 930 |
| Work objects. . .  | 931 |
| Example: Work object . . .                               | 932 |
| Developing event listeners . . .                         | 933 |
| Using the application notification service . . .         | 933 |
| Example: Event listener . . .                            | 934 |
| Developing Asynchronous scopes . . .                     | 935 |
| Asynchronous scopes. . .                                 | 937 |
| Alarms . . .   | 938 |
| Subsystem monitors . . .                                 | 938 |
| Asynchronous scopes: Dynamic message bean scenario . . . | 939 |
| Interoperating with asynchronous beans . . .             | 940 |
| Asynchronous beans interoperability issues . . .         | 940 |

## **Chapter 18. Using object pools. . . . . 943**

|  |     |
|--|-----|
| Object pool managers . . .                   | 944 |
| Object pool manager collection . . .         | 946 |
| Name . . .                                   | 946 |
| JNDI Name . . .                              | 946 |
| Description . . .                            | 946 |
| Category . . .                               | 946 |
| Object pool manager settings . . .           | 946 |
| Object pool service settings. . .            | 947 |
| Startup . . .                                | 947 |
| Object pools: Resources for learning . . .   | 948 |
| Object pool performance considerations . . . | 949 |

## **Chapter 19. Using startup beans . . . . . 951**

## Chapter 20. Using the scheduler service . . . . . 953

|   |     |
|---|-----|
| Managing the scheduler service . . . . .              | 953 |
| Creating the database for scheduler . . . . .         | 953 |
| Configuring a scheduler . . . . .                     | 958 |
| Enabling the scheduler service. . . . .               | 962 |
| Developing and scheduling tasks . . . . .             | 963 |
| Developing a task that calls a session bean . . . . . | 963 |
| Developing a task that sends a JMS message . . . . .  | 964 |
| Receiving scheduler notifications . . . . .           | 965 |
| Submitting a task to a scheduler . . . . .            | 966 |
| Managing tasks with a scheduler . . . . .             | 967 |
| Scheduler interface . . . . .                         | 969 |
| Interoperating with the Scheduler service . . . . .   | 972 |
| Recreating Scheduler tasks . . . . .                  | 973 |
| Deleting Scheduler tasks. . . . .                     | 974 |
| Recreating Scheduler tables. . . . .                  | 975 |

## Chapter 21. Using shared work areas 977

|   |     |
|---|-----|
| WorkArea service - Overview . . . . .                 | 977 |
| Work area property modes . . . . .                    | 978 |
| Nested work areas . . . . .                           | 979 |
| Distributed work areas . . . . .                      | 981 |
| WorkArea service: Special considerations . . . . .    | 981 |
| WorkArea service performance considerations . . . . . | 982 |
| Developing applications that use work areas . . . . . | 983 |
| UserWorkArea interface . . . . .                      | 983 |
| Example: WorkArea SimpleSample application . . . . .  | 984 |
| Accessing the WorkArea service . . . . .              | 985 |
| Beginning a new work area . . . . .                   | 985 |
| Setting properties in a work area . . . . .           | 986 |
| Using a work area to manage local work . . . . .      | 987 |
| Completing a work area. . . . .                       | 991 |
| Managing the work area service . . . . .              | 992 |
| Enabling the WorkArea service . . . . .               | 992 |
| Managing the size of work areas . . . . .             | 994 |

## Chapter 22. Using the transaction service . . . . . 995

|   |      |
|---|------|
| Transaction support in WebSphere Application Server . . . . .                 | 995  |
| Resource manager local transaction (RMLT) . . . . .                           | 996  |
| Global transactions . . . . .   | 997  |
| Local transaction containment (LTC). . . . .                                  | 997  |
| Local and global transaction considerations . . . . .                         | 1001 |
| Developing components to use transactions . . . . .                           | 1002 |
| Setting transactional attributes in the deployment descriptor . . . . .       | 1002 |
| Using bean-managed transactions . . . . .                                     | 1004 |
| Configuring transaction properties for an application server . . . . .        | 1005 |
| Transaction service settings . . . . .  | 1006 |
| Managing active transactions. . . . .   | 1007 |
| Managing transaction logging for optimum server availability . . . . .        | 1008 |
| Configuring transaction aspects of servers for optimum availability . . . . . | 1009 |
| Moving a transaction log from one server to another . . . . .                 | 1011 |

|   |      |
|---|------|
| Restarting an application server on a different host . . . . .  | 1012 |
| Transactional interoperation with non-WebSphere application servers . . . . .   | 1012 |
| Troubleshooting transactions . . . . .  | 1012 |
| Transaction service exceptions . . . . .  | 1013 |
| UserTransaction interface - methods available . . . . .   | 1014 |
| Coordinating access to 1-PC and 2-PC-capable resources within the same transaction . . . . .                          | 1014 |
| Coordinating access to 1-PC and 2-PC-capable resources within the same transaction . . . . .                          | 1015 |
| Enabling an application to coordinate access to 1-PC and 2-PC-capable resources within the same transaction . . . . . | 1016 |
| Configuring an application server to allow logging for heuristic reporting . . . . .                                  | 1017 |
| Exceptions thrown for transactions involving both single- and two-phase commit resources . . . . .                    | 1017 |
| Last Participant Support: Resources for learning . . . . .  | 1017 |

## Chapter 23. Using the ActivitySession service . . . . . 1019

|   |      |
|---|------|
| The ActivitySession service . . . . .   | 1020 |
| Using ActivitySessions with HTTP sessions . . . . .                                 | 1021 |
| ActivitySession and transaction contexts . . . . .                                  | 1023 |
| Combining transaction and ActivitySession container policies. . . . .               | 1024 |
| Developing a J2EE application to use ActivitySessions . . . . .                     | 1030 |
| Developing an enterprise bean or J2EE client to manage ActivitySessions . . . . .   | 1032 |
| Configuring ActivitySession deployment attributes for an enterprise bean . . . . .  | 1033 |
| Container ActivitySession assembly properties for EJB modules . . . . .             | 1035 |
| Configuring ActivitySession deployment attributes for a Web application . . . . .   | 1037 |
| Disabling or enabling the ActivitySession service . . . . .                         | 1038 |
| ActivitySession service settings . . . . .  | 1039 |
| Configuring the default ActivitySession timeout for an application server . . . . . | 1040 |
| ActivitySession service settings . . . . .  | 1041 |
| Troubleshooting ActivitySessions . . . . .  | 1042 |
| The ActivitySession service application programming interfaces . . . . .            | 1042 |
| Samples: ActivitySessions . . . . .   | 1043 |
| ActivitySession service: Resources for learning . . . . .                           | 1044 |

## Chapter 24. Using naming . . . . . 1047

|   |      |
|---|------|
| Naming . . . . .  | 1048 |
| New features for name space support . . . . .                             | 1048 |
| Name space logical view . . . . .   | 1049 |
| Name space partitions . . . . .   | 1050 |
| Initial context support . . . . .   | 1052 |
| Initial contexts registered with the ORB as initial references . . . . .  | 1052 |
| Default initial contexts . . . . .  | 1053 |
| Lookup names support in deployment descriptors and thin clients . . . . . | 1053 |

|  |      |
|--|------|
| Relative names . . . . .   | 1053 |
| Qualified names . . . . .  | 1054 |
| JNDI support in WebSphere Application Server   | 1056 |
| Developing applications that use JNDI . . . . .  | 1056 |
| Example: Getting the default initial context   | 1058 |
| Example: Getting an initial context by setting<br>the provider URL property . . . . .                                | 1062 |
| Example: Setting the provider URL property to<br>select a different root context as the initial<br>context . . . . . | 1064 |
| Example: Looking up an EJB home with JNDI  | 1065 |
| Example: Looking up a JavaMail session with<br>JNDI . . . . .  | 1067 |
| JNDI interoperability considerations . . . . .   | 1068 |
| JNDI caching . . . . .   | 1069 |
| JNDI cache settings . . . . .  | 1070 |
| Example: Controlling JNDI cache behavior<br>from a program . . . . .   | 1072 |
| JNDI name syntax . . . . .   | 1073 |
| INS name syntax . . . . .  | 1073 |
| JNDI to CORBA name mapping considerations  | 1073 |
| Example: Setting the syntax used to parse<br>name strings . . . . .  | 1074 |
| Developing applications that use CosNaming<br>(CORBA Naming interface) . . . . .                                     | 1074 |
| Example: Getting an initial context with<br>CosNaming . . . . .  | 1074 |
| Example: Looking up an EJB home with<br>CosNaming . . . . .  | 1077 |
| Configured name bindings . . . . .   | 1079 |
| Configured binding types . . . . .   | 1080 |
| Name space federation . . . . .  | 1081 |
| Name space bindings . . . . .  | 1083 |
| Configuring and viewing name space bindings  | 1083 |
| String binding settings . . . . .  | 1083 |
| CORBA object binding settings . . . . .  | 1084 |
| Indirect lookup binding settings. . . . .  | 1085 |
| EJB binding settings. . . . .  | 1085 |
| Name space binding collection . . . . .  | 1086 |
| Configuring name servers . . . . .   | 1086 |
| Name server settings . . . . .   | 1087 |
| Troubleshooting name space problems . . . . .  | 1087 |
| dumpNameSpace tool . . . . .   | 1087 |
| Example: Invoking the name space dump<br>utility . . . . .   | 1090 |
| Name space dump utility for java:, local:and<br>server name spaces . . . . .   | 1090 |
| Name space dump sample output . . . . .  | 1092 |
| Naming and directories: Resources for learning   | 1094 |

## Chapter 25. Using the dynamic cache service to improve performance. . . 1097

|  |      |
|--|------|
| Dynamic cache . . . . .                            | 1097 |
| Configuring cache replication . . . . .            | 1098 |
| Cache replication. . . . .                         | 1098 |
| Internal messaging configuration settings. . . . . | 1099 |
| Cache replication. . . . .                         | 1100 |
| Enabling the dynamic cache service . . . . .       | 1100 |
| Dynamic cache service settings . . . . .           | 1100 |
| Configuring servlet caching . . . . .              | 1101 |
| Configuring cache replication. . . . .             | 1102 |

|  |      |
|--|------|
| Configuring the dynamic cache disk offload                             | 1104 |
| Configuring Edge Side Include caching . . . . .                        | 1105 |
| Configuring external cache groups . . . . .                            | 1107 |
| Displaying cache information. . . . .                                  | 1112 |
| Configuring cacheable objects with the<br>cachespec.xml file . . . . . | 1112 |
| Verifying the cacheable page . . . . .                                 | 1114 |
| Cachespec.xml file . . . . .   | 1114 |
| Configuring command caching . . . . .                                  | 1120 |
| Command class . . . . .  | 1121 |
| CacheableCommandImpl class . . . . .                                   | 1121 |
| Example: Caching a command object . . . . .                            | 1122 |
| Using the DistributedMap interface for the<br>dynamic cache . . . . .  | 1123 |
| Sharing cached objects in a clustered<br>environment . . . . .         | 1125 |
| Cache instance settings. . . . .                                       | 1126 |
| Cache instance collection . . . . .                                    | 1127 |
| Cache instance service settings . . . . .                              | 1127 |
| Invalidation listeners . . . . .                                       | 1128 |
| Example: Caching Web services . . . . .                                | 1128 |
| Example: Configuring the dynamic cache. . . . .                        | 1131 |
| Cache monitor. . . . .   | 1133 |
| Edge cache statistics. . . . .   | 1134 |

## Chapter 26. Managing user profiles 1137

|  |      |
|--|------|
| User profile . . . . .   | 1137 |
| UserProfileManager class . . . . .   | 1138 |
| User profile development options . . . . .   | 1138 |
| Extending the data represented in user profiles  | 1138 |
| Adding columns to the base user profile<br>implementation . . . . .                    | 1138 |
| Extending the User Profile enterprise bean and<br>importing legacy databases . . . . . | 1139 |
| UPServletExample.java . . . . .  | 1139 |
| UserProfileExtendedSample.java. . . . .  | 1141 |
| UPServletExampleExtended.java. . . . .   | 1142 |
| UserProfileExtended.java . . . . .   | 1144 |
| UPServletExtended.java . . . . .   | 1145 |
| userprofile.xml . . . . .  | 1147 |

## Chapter 27. Assembling applications 1149

|  |      |
|--|------|
| Application assembly and J2EE applications. . . . .                  | 1150 |
| Archive support in Version 5.0 . . . . .                             | 1151 |
| Starting the Application Assembly Tool (AAT)                         | 1151 |
| Migrating application modules from J2EE 1.2 to<br>J2EE 1.3 . . . . . | 1152 |
| earconvert tool . . . . .  | 1153 |
| Assembling new or modifying existing modules                         | 1154 |
| Adding files to assembled modules. . . . .                           | 1156 |
| Resource environment reference assembly<br>settings . . . . .        | 1157 |
| Resource Adapter Archive file assembly<br>settings . . . . .         | 1158 |
| Saving applications after assembly . . . . .                         | 1161 |
| Verifying archive files . . . . .                                    | 1162 |
| Application assembly performance checklist. . . . .                  | 1163 |
| Generating code for deployment . . . . .                             | 1163 |
| ejbdeploy tool . . . . .   | 1164 |

|   |      |
|---|------|
| ejbdeploy syntax — relationship to Application    |      |
| Assembly Tool options . . . . .                   | 1164 |
| Application Assembly Tool: Resources for learning | 1165 |

**Chapter 28. Deploying and managing applications . . . . . 1167**

|   |      |
|---|------|
| Enterprise applications . . . . .               | 1167 |
| Installing a new application . . . . .          | 1167 |
| Preparing for application installation settings | 1172 |
| Example: Installing an EAR file using the       |      |
| default bindings . . . . .                      | 1176 |
| Enterprise application collection . . . . .     | 1176 |

|  |      |
|--|------|
| Name . . . . .                                 | 1176 |
| Status . . . . .                               | 1177 |
| Enterprise application settings . . . . .      | 1177 |
| Starting and stopping applications . . . . .   | 1179 |
| Exporting applications . . . . .               | 1180 |
| Exporting DDL files . . . . .                  | 1180 |
| Updating applications . . . . .                | 1180 |
| Hot deployment and dynamic reloading . . . . . | 1182 |
| Uninstalling applications . . . . .            | 1190 |
| Deploying and managing applications: Resources |      |
| for learning. . . . .                          | 1191 |

---

## Trademarks and service marks

The following terms are trademarks of IBM Corporation in the United States, other countries, or both:

- AIX
- CICS
- Cloudscape
- DB2
- Everyplace
- iSeries
- IBM
- Informix
- iSeries
- MQSeries
- OS/390
- Redbooks
- SupportPac
- ViaVoice
- VisualAge
- WebSphere
- zSeries

Java and all Java-based trademarks are trademarks of Sun Microsystems, Inc. in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Other company, product and service names may be trademarks or service marks of others.



---

## Chapter 1. Welcome to Applications

The following items comprise the application programming model, including numerous services available to support deployed applications.

### Web modules

Use Web components such as servlets and JavaServer Pages files to develop dynamic Web sites. Product extensions to the open source servlet and JSP APIs enhance standard features, and provide additional functionality.

Web modules consist of the following application components, each performing a different function:

- HTML and JSP pages provide the user interface and program logic
- Servlets coordinate work between other components of the application
- Bean Scripting Framework enables users to implement JavaScript code in JSP applications

HTTP sessions are a key area of product support for Web modules. By managing **0 sessions** for your Web applications, you can personalize a Web site for individual customers. A session is a series of requests to a servlet, originating from the same user at the same browser. Managing HTTP sessions allows servlets running in a Web container to keep track of individual users. For example, a servlet might use sessions to provide "shopping carts" to on-line shoppers. Suppose the servlet is designed to record the items each shopper indicates he or she will purchase from the Web site. It is important that the servlet be able to associate incoming requests with particular shoppers. Otherwise, the servlet might mistakenly add choices of Shopper 1 to the cart of Shopper 2.

### EJB modules

IBM WebSphere Application Server provides broad support for enterprise beans, including the Enterprise JavaBeans (EJB) 2.0 specification. The EJB 2.0 specification introduces a container-managed persistence (CMP) 2.0 component model, which provides a number of improvements to aid developer productivity and application performance. In addition, this product continues to fully support enterprise beans written to the CMP 1.1 programming model and deployed in previous versions of this product; applications can use CMP 1.1 beans, CMP 2.0 beans, or a mixture of both. CMP 1.1 beans can be directly carried forward in an EJB 1.1 ejb-jar module or may be repackaged and combined with CMP 2.0 beans in an EJB 2.0 module.

For EJB 2.0 modules, a feature introduced in Version 5 of this product, called **access intent** policies, eases the management of interactions between CMP beans and their underlying datastores. Each policy sets such data access characteristics such as access type (read or update) and transaction isolation that affect the locking of resources, letting you choose the level of data integrity and performance for your application. The Enterprise Extensions product adds APIs to enable you to further customize IBM-provided access intent policies for your particular environment.

Several excellent trade books that cover EJB 2.0 and the CMP 2.0 persistence model are already available. A good way to locate some of these is to visit your favorite

online bookstore and search on the term *Enterprise JavaBeans*. For a more basic orientation, see ("Enterprise beans: Resources for learning").

Your application development might include **asynchronous messaging**, which the product supports as a method of communication based on the Java Message Service (JMS) programming interface.

The base JMS support enables IBM WebSphere Application Server applications to exchange messages asynchronously with other JMS clients by using JMS destinations (queues or topics). An application can explicitly poll for messages on a destination.

The product also provides a message listener service that applications can use to automatically retrieve messages from JMS destinations for processing by message-driven beans, without the application having to explicitly poll JMS destinations.

The Enterprise Extensions product provides *extended messaging*, which uses the EJB container to manage the messaging infrastructure, and provides more types of messaging beans. This enables application developers to concentrate on the business logic for enterprise beans and to leave the messaging usage to messaging objects and configuration of the EJB container.

Referto:

- "Asynchronous messaging with WebSphere - an overview" (not in this document)An overview of WebSphere asynchronous messaging.
- "Using JMS and messaging in applications" (not in this document)Implementing WebSphere J2EE applications that use JMS.

### **Client modules**

The Imaginary Buffer Line product provides a CD-ROM and installation program for installing application clients without installing the entire application server. This smaller footprint is useful if you want to run client applications on multiple client machines. For more information, see ("Installing application clients").

Application clients follow several programming models with unique requirements and suitability for different types of applications. Most of these models are only available when you install the product. Models include:

#### **ActiveX application client**

Supported Windows platforms only; Client only

#### **Applet client**

Supported Windows platforms only; Client only

#### **J2EE application client**

All supported server platforms; All supported client platforms

#### **Pluggable application client**

Supported Windows platforms only, Sun JRE 1.3.1\_03 or later (but not 1.4); Thin application client programming model only (no J2EE programming model); Client only

#### **Thin application client**

All supported client platforms

### **Web services**



The Web services components included with this product version build upon the Apache Simple Object Access Protocol (SOAP) 2.3-based capabilities delivered with Version 4.0.x of the product.

Version 5.0 introduced an open source implementation for a Web Services Invocation Framework (WSIF).

An additional Web services component, "IBM WebSphere Web Services for J2EE Technology Preview" is available for use with Version 5. It is available as a separate download from:  
<http://www7b.boulder.ibm.com/wsdd/downloads/techpreviews.html>.

The Web Services Technology Preview supports emerging Java Web services standards like JAX-RPC and Web services for J2EE. It is recommended that new development efforts use the Web Services Technology Preview and follow these standards.

Additional features, such as UDDI Registry, and Web Services Gateway are described in "Welcome to Servers" (not in this document).

## **5.0.2**

WebSphere Application Server supports Web services security functionality that is based on standards included in the Web services security (WS-Security) specification.

### **CORBA applications**

The Enterprise Extensions product enables the creation of CORBA client and server applications within a IBM WebSphere Application Server environment. You can use the CORBA C++ SDK to build a lightweight WebSphere CORBA server to use with new or existing C and C++ programs. You also can use the SDK to build a WebSphere CORBA C++ client to use with a WebSphere EJB server or WebSphere CORBA C++ server. See ("Implementing CORBA applications")Implementing CORBA applications

### **Application services**

IBM WebSphere Application Server provides essential services to ease the building of dynamic and flexible e-business applications. These services support and extend the open standards of J2EE and Web services, with a focus on application reuse and integration.

The Enterprise Extensions product takes application services to the next level, providing a broad range of dynamic API extensions that address functional gaps in the J2EE programming model.

- **Class loading**

The WebSphere Application Server product provides several class-loading modes, policies, and features to enable you to deploy and run your applications successfully. An application server provides an Application class-loader policy that enables you to control the isolation of applications in a server. If you want applications to share classes, choose the SINGLE policy; otherwise choose the MULTIPLE policy, which isolates the class loaders for each application.

Similarly, at the application level, you can choose a (WAR class-loader policy) that configures the isolation of Web modules within an application. If you

choose the policy APPLICATION, then each Web module in your application can see the classes of other Web modules. A policy of MODULE creates a separate class loader for each Web module, resulting in isolation for each the classes of each Web module.

The class-loader mode setting, which you can configure at the server, application, or Web module level depending on your class-loader policy, enables you to control whether application class loaders override classes contained in base run-time class loaders. By default, the WebSphere Application Server class loaders have a class-loader mode of PARENT\_FIRST, which is the standard JDK mode and does not allow the application class loader to override classes. You must take care when using the PARENT\_LAST class-loader mode to make all dependent classes available within the application or you might get LinkageErrors or other class-loader exceptions. For example, if you provide a newer version of the Xerces.jar file and your application is using XSLT, you must also provide a xalan.jar file within your application.

- **Shared library**

Version 5.0 of WebSphere Application Server introduces the concept of a "Managing shared libraries" (not in this document). A shared library is a CLASSPATH and a symbolic name for that classpath. You define shared libraries at the cell, node, or server level and then associate the shared libraries either with an application server (making the classes available to all applications in the server) or with individual applications (making the classes available only to the referencing application). This mechanism provides a convenient way to make libraries of classes available to your applications outside of a standard J2EE enterprise application (EAR) file for easier version management and space efficiency.

- **EJB query**

The EJB query language is used to specify a query over container-managed entity beans. WebSphere's EJB query language is compliant with the EJB QL defined in Sun's EJB 2.0 specification, but adds additional support as described in the topic ("Comparison of EJB 2.0 specification and WebSphere query language").

EJB query can be used to define a finder or select method of an EJB entity bean. Finder and select queries are specified in the bean deployment descriptor using the <ejb-ql> tag. Queries specified in the deployment descriptor are compiled into SQL during deployment. See ("Using EJB query").

The Enterprise Extensions product includes the dynamic query service, an additional API that enables you to dynamically specify a query in your application by adding the executeQuery() method. See ("Using the dynamic query service").

- **Internationalization support**

The internationalization service manages the distribution of locale and time zone information, or *internationalization context*, in applications that run on WebSphere Application Server Enterprise installations. The internationalization service solves the problem of mismatched locales and time zones by systematically managing the distribution of internationalization context across the various components of EJB applications.

### 5.0.2

The internalization service transparently propagates internationalization context over requests that originate from J2EE-compliant Web service clients. The service creates a SOAP header block that contains the invocation context scoped to the current thread; this SOAP representation is then inserted into the outgoing Web

services request. For incoming requests, the service scopes the propagated internationalization context, referred to as *caller context*, to the invocation of the stateless session bean that is enabled as a Web service. The service also scopes an *invocation context* as prescribed by the internationalization context management policies that were assigned to the enterprise bean's methods during application assembly.

For more information about the internationalization service, see ("Using the internationalization service")

- **Workarea service**

The WorkArea service enables application developers to implicitly propagate information beyond the information passed in remote calls. Applications can create a work area, insert information into it, and make remote invocations. The work area is propagated with each remote method invocation, eliminating the need to explicitly include an appropriate argument in the definition of each method. The methods on the server side can use or ignore the information in the work area as appropriate. See ("Using shared work areas")

- **Application profiling**

Application profiling enables you to configure multiple access intent policies on the same method of an entity bean; and to configure multiple access intent policies for dynamic query on the same entity bean.

To use application profiling, application developers identify named units of work, or *tasks*. A task typically corresponds to the execution of a concrete and high-level job within the application. The IBM WebSphere Application Server run-time environment queries the task at the invocation of any entity bean, and establishes the appropriate access intent policy under which the bean should execute. An *application profile* is the set of access intent or query intent policies that should be selectively applied, as well as the list of tasks for which the policies should be applied. See ("Application profiling")

- **Scheduler service**

The scheduler service enables J2EE work to be executed at a requested time or interval. The scheduler API supports different implementations of the TaskInfo interface, each of which can be used to schedule a particular type of work; for example, you can develop a task that calls a session bean or a task that sends a JMS message. You can set a notification sink on a task in order to receive the notification events that are generated by a scheduler when it performs an operation on the task. See ("Using the scheduler service")

- **Asynchronous beans**

An asynchronous bean is a Java object or enterprise bean that can be executed asynchronously by a J2EE application, using the J2EE context of the bean's creator. These beans also can run with copies of other J2EE contexts. For example:

- Internationalization context
- Application profiles
- Work areas
- Access intent policies

Asynchronous beans enable the construction of stateful, "active" J2EE applications. These applications address a segment of the application space that J2EE has not previously addressed (that is, advanced applications that require application threading, active agents within a server application, or distributed monitoring capabilities). See ("Using asynchronous beans").

- **Object pools**

Objects are frequently pooled by Java applications in order to avoid the cost of creating new Java objects and the associated garbage collection delays that result when these objects are reclaimed after use. An object pool keeps a number of pre-allocated objects on behalf of its users. Applications can get an object from the pool, use it, and later return it to the pool. This allows the individual object instances to be reused and effectively limits the amount of garbage generated by the application. ("Using object pools")

- **Startup beans**

Startup beans are stateful session beans that enable J2EE applications to execute business logic when an application starts or stops. The startup bean is loaded when the application starts. The start() method is then invoked on the bean's remote interface. This method can execute any business logic needed by the application at start time. Similarly, the bean's stop() method is called on the instance when the application is stopped and can execute any business logic needed by the application at stop time. See ("Using startup beans")

Startup beans are especially useful when used in combination with asynchronous beans to develop an active J2EE server application.

- **Business Rule Beans (BRBeans)**

Business Rule Beans are used to separate business rules from an application's core behavior, allowing the application code to remain intact and untouched even as business practices change. Each business rule is represented by an entity bean that persistently stores information related to that rule. Each business rule is assigned an appropriate rule name and stored in a rule folder. The application developer identifies "points of variability" within an application and codes trigger points at these locations. These trigger points invoke one or more business rules. See ("Using Business Rule Beans")

- **Transactions**

IBM WebSphere Application Server applications can use transactions to coordinate multiple updates to resources as atomic units (as indivisible units of work) such that all or none of the updates are made permanent. The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself)
- Entity beans use container-managed transactions
- Web components (servlets) use bean-managed transactions

The product is a transaction manager that supports the coordination of resource managers through their XAResource interface and participates in distributed global transactions with other OTS 1.2 compliant transaction managers (for example J2EE 1.3 application servers). Applications can also be configured to interact with databases, JMS queues, and JCA connectors through their local transaction support when distributed transaction coordination is not required.

Resource managers that offer transaction support can be categorized into those that support 2-phase coordination (by offering an XAResource interface) and those that support only 1-phase coordination (for example through a LocalTransaction interface). The IBM WebSphere Application Server transaction support provides coordination, within a transaction, for any number of 2-phase capable resource managers. It also enables a single 1-phase capable resource manager to be used within a transaction in the absence of any other resource

managers, although a WebSphere transaction is not necessary in this case. With the Last Participant Support of Enterprise Extensions, you can coordinate the use of a single 1-phase commit (1PC) capable resource with any number of 2-phase commit (2PC) capable resources in the same global transaction. At transaction commit, the 2-phase commit resources are prepared first using the 2-phase commit protocol, and if this is successful the 1-phase commit-resource is then called to commit(one\_phase). The 2-phase commit resources are then committed or rolled back depending on the response of the 1-phase commit resource.

The ActivitySession service of Enterprise Extensions provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. It is a distributed context that can be used to coordinate multiple 1-phase resource managers. The product EJB container and deployment tooling support ActivitySessions as an extension to the J2EE programming model. Enterprise beans can be deployed with lifecycles that are influenced by ActivitySession context, as an alternative to transaction context. An application can then interact with a resource manager through its LocalTransaction interface for the period of a client-scoped ActivitySession rather than just the duration of an EJB method.

- **Naming**

Naming clients use (**Naming Services**) primarily to access objects, such as EJB homes, associated with applications installed on IBM WebSphere Application Server. Objects are made available to clients by being bound into a name space. A name space is under the control of a name server. In this product, there are potentially many name servers, and the name spaces controlled by the various name servers are federated together to form the view of a single name space. Each name server presents the same logical view of the federated name spaces.

Name servers provided by this product are a CORBA CosNaming implementation. IBM WebSphere Application Server provides a CosNaming JNDI plug-in which enables clients to access the name servers through the JNDI interface. Clients to EJB applications typically use JNDI to perform Naming operations. Clients may access the name servers directly through the CORBA programming model. The CosNaming interface is part of the CORBA programming model. CORBA clients which need to access EJB homes or some other objects bound to the name space would typically use the CORBA CosNaming interface to perform Naming operations.

- **Dynamic cache**

Dynamic cache improves application performance by caching outputs and contents of outputs of Servlets, JavaServer Pages (JSP) files, Web Services, and Commands. On subsequent client requests to the same applications, Dynamic cache intercepts these calls and responds by serving the output or the contents of output from the cache.

Dynamic cache in this product version includes:

- **Servlet/JSP files caching**

This caches output of dynamic servlets and JSP files by working with Java virtual machine of the application server by intercepting calls to service methods and serving Web pages from the cache. This improves server response time, throughput and scalability.

- **Command caching**

Commands that are written to the Command Architecture encapsulate business logic tasks and provide a standard way to invoke the business logic request. Command objects need to implement CacheableCommand interface instead of TargetableCommand interface to cache. Like in servlets and JSP caching, requests to execute business logic in the

command is intercepted by the cache. If a command with the same request attributes are available in cache, output properties are copied from the cached instance to the requested instance and returned without executing the business logic again.

### **Web Services caching**

Web service responses can be cached just like servlet and JSP results. These requests are intercepted and cache ID computed based on how the cache ID rules are specified in the cache policy. Hash of the whole SOAPEnvelope can be used as a cache ID or it can be parsed and service, operation and parameters to these operations used as cache ID. If a cache entry is not found for the computed cache id, the request is forwarded to the SOAP engine and the result is cached.

### **Edge Side Include caching**

This provides the ability to cache, assemble and deliver dynamic web pages at the edge of the enterprise network. Edge Side Includes (ESI) is a simple markup language which enables dynamic web pages (which by themselves are not so cache efficient) to be broken down into cacheable fragments. These fragments are then cached on the edge of the network and assembled into a single page upon user requests.

### **Distributed caching**

Cache contents can be shared and replicated among servers by Dynamic Caching using an underlying JMS based message broker system, DRS (Data Replication Service). Sharing characteristics of individual cache entry is configured using the cache policy specification.

- **User profiles**

Managing (**user profiles**) allows a company to maintain database tables containing fields for demographic data of individual customers or other users on the company system. For example, when a user repeatedly logs onto a Web site that supports user profiles, the Web site can display headlines and advertising tailored to the shopping preferences of that user. The site can address the user by his or her logon name. User profile API is deprecated in the current release.

### **Service choreography**

The Enterprise Extensions product introduces service choreography, a feature that enables you to plan and implement complex business processes (workflows). You can use the Process Choreographer to model task-oriented business services as *processes*. Refer to: ("Using Process Choreographer").

### **Assembly tools**

The Application Assembly Tool and a command line deployment tool are provided for packaging your application code components into the needed modules for deployment onto the server.

See ("Assembling applications").

EAR files are comprised of the following archives:

- Enterprise bean (JAR) files (known as ("EJB modules"))
- Web application (WAR) files (known as ("Web modules"))
- Application client (JAR) files (known as ("Application clients"))
- Resource adapter (RAR) files (known as resource adapter modules)

- Optionally, additional JAR files containing dependent classes or other components required by the application

The standard file extension of an Enterprise application file is .ear.

For a discussion of archives and Web components supported by the Application Assembly Tool in Version 5, see ("Archive support in Version 5.0").

See the "Assembling or packaging" (not in this document) topic for more information.

### **5.0.2**

For the Windows and Linux Intel operating systems, the Assembly Toolkit replaces the Application Assembly Tool (AAT). Visit [http://www-3.ibm.com/software/webservers/appserv/was/support/\(http://www-3.ibm.com/software/webservers/appserv/was/support/\)](http://www-3.ibm.com/software/webservers/appserv/was/support/(http://www-3.ibm.com/software/webservers/appserv/was/support/)) to download the Assembly Toolkit, which offers the Assembly Toolkit and other products. The Assembly Toolkit consists of the J2EE Perspective of the WebSphere Studio Application Developer product, without the code generation capabilities.

## **Deployment**

Tools, such as the WebSphere Administrative Console, are provided for installing your modules onto the application server, then managing the installed modules in their respective containers.





---

## Chapter 2. Using Web applications

A developer creates the files comprising a Web application, and then assembles the Web application components into a Web module. Next, the deployer (typically the developer in a unit-testing environment or the administrator in a production environment) installs the Web application on the server.

### Steps for this task

1. **(Optional)** Migrate existing Web applications to run in the new version of WebSphere.
2. Design the Web application and develop its code artifacts: Servlets, JavaServer Pages (JSP) files, and static files, as for example, images and Hyper Text Markup Language (HTML) files.

See the "Resources for learning" article for links to design documentation.

3. **(Optional)** Implement JavaScript within JSP tags using the Bean Scripting Framework (BSF).
4. Develop the Web application, using WebSphere Application Server extensions to enhance its functionality.
5. Assemble the Web application into a Web module.

Web module assembly properties might include the ability to:

- Configure servlet page lists
- Configure servlet filters
- Serve servlets by class name
- Enable file serving

6. Deploy the Web module or application module that contains the Web application.

Following deployment, you might find it handy to use the tool that enables batch compiling of the JSP files for quicker initial response times.

7. **(Optional)** Troubleshoot your Web application. (Refer to "Web container troubleshooting tips" in WebSphere Application Server documentation on monitoring and troubleshooting.)
8. **(Optional)** Modify the default Web container configuration in the application server in which you deployed the Web module or application module containing the Web application.
9. **(Optional)** Manage the deployed Web application..

---

## Web applications

A Web application is comprised of one or more related servlets, JavaServer Pages technology (JSP files), and Hyper Text Markup Language (HTML) files that you can manage as a unit.

The files in a Web application are related in that they work together to perform a business logic function.

For example, one of the WebSphere Application Server samples is a Simple Greeting Web application. This application, comprised of a servlet and Web pages, greets new users when the application is accessed.

The Web application is a concept supported by the Java Servlet Specification. Web applications are typically packaged as .war files.

---

## web.xml file

The web.xml file provides configuration and deployment information for the Web components that comprise a Web application. Examples of Web components are servlet parameters, servlet and JavaServer Pages (JSP) definitions, and Uniform Resource Locators (URL) mappings.

The servlet 2.3 specification dictates the format of the web.xml file, which makes this file portable among Java Two Enterprise Edition (J2EE) compliant products.

### Location

The web.xml file must reside in the WEB-INF directory under the context of the hierarchy of directories that exist for a Web application. For example, if the application is client.war, then the web.xml file is placed in the *install\_root/client.war*/WEB-INF directory.

### Usage notes

- Is this file read-only?  
No
- Is this file updated by a product component?  
This file is updated by the Application Assembly Tool (AAT).
- If so, what triggers its update?  
The AAT updates the web.xml file when you assemble Web components into a Web module, or when you modify the properties of the Web components or the Web module.
- How and when are the contents of this file used?  
WebSphere Application Server functions use information in this file during the configuration and deployment phases of Web application development.

### Sample file entry

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE web-app PUBLIC "-//Sun Microsystems, Inc.//DTD Web Application 2.3//EN"
"http://java.sun.com/j2ee/dtds/web-app_2_3.dtd">

  <web-app id="WebApp_1">
    <display-name>Persistence Manager Web Client</display-name>
    <description>Peristence Manager Web Client</description>
    <servlet id="Servlet_1">
      <servlet-name>CustomerLocalServlet</servlet-name>
      <description>Local Customer Servlet</description>
      <servlet-class>CustomerLocalServlet</servlet-class>
    </servlet>
    <servlet id="Servlet_2">
      <servlet-name>CustomerServlet</servlet-name>
      <description>Remote Customer Servlet</description>
      <servlet-class>CustomerServlet</servlet-class>
    </servlet>
    <servlet id="Servlet_3">
      <servlet-name>CreditCardServlet</servlet-name>
      <description>Credit Card Servlet - PM Verification</description>
      <servlet-class>CreditCardServlet</servlet-class>
    </servlet>
    <servlet-mapping id="ServletMapping_1">
      <servlet-name>CustomerLocalServlet</servlet-name>
```

```

        <url-pattern>/CustomerLocal</url-pattern>
    </servlet-mapping>
    <servlet-mapping id="ServletMapping_2">
        <servlet-name>CustomerServlet</servlet-name>
        <url-pattern>/Customer</url-pattern>
    </servlet-mapping>
    <servlet-mapping id="ServletMapping_3">
        <servlet-name>CreditCardServlet</servlet-name>
        <url-pattern>/CreditCard</url-pattern>
    </servlet-mapping>
    <welcome-file-list id="WelcomeFileList_1">
        <welcome-file>index.html</welcome-file>
    </welcome-file-list>
    <security-role id="SecurityRole_1">
        <description>Everyone role</description>
        <role-name>Everyone Role</role-name>
    </security-role>
    <security-role id="SecurityRole_2">
        <description>AllAuthenticated role</description>
        <role-name>All Role</role-name>
    </security-role>
    <security-role id="SecurityRole_3">
        <description>Deny all access role</description>
        <role-name>DenyAllRole</role-name>
    </security-role>
</web-app>

```

---

## Migrating Web application components

Supported open specification levels in WebSphere Application Server Version 5 are documented in article, "Migrating" (not in this document).

Migration of Web applications deployed in WebSphere Application Server Version 4.x is not necessary; version 2.2 of the servlet specification and version 1.1 of the JavaServerPages (JSP) specification are still supported. However, where there are behavioral differences between the Java Two Enterprise Edition (J2EE) 1.2 and J2EE 1.3 specifications, bear in mind that J2EE 1.3 specifications are implemented in WebSphere Application Server Version 5 and will override any J2EE 1.2 behaviors.

Servlet migration might be a concern if your application:

- implements a WebSphere internal servlet to bypass a WebSphere Application Server Version 4.x single application path restriction.
- extends a PageListServlet that relies on configuration information in the servlet configuration XML file.
- uses a servlet to generate Hyper Text Markup Language (HTML) output.
- calls the `response.sendRedirect()` method for a servlet using the `encodeRedirectURL` function or executing within a non-context root.

JSP migration might be a concern if your application references JSP page implementation classes in unnamed packages, or if you install WebSphere Application Server Version 4.x EAR files (deployed in Version 4.x with the JSP Precompile option), in Version 5.

Follow these steps if migration issues apply to your Web application:

### Steps for this task

1. Use WebSphere Application Server Version 5 package names for any WebSphere Application Server Version 4.x internal servlets, which are implemented in your application.

In WebSphere Application Server Version 4.x, Web modules with a context root setting of / are not supported. Accessing Web modules with this root context results in HTTP 404 - File not Found errors.

To bypass the errors, and to enable the serving of static files from the root context, WebSphere Application Server Version 4.x users are advised to add the servlet class, `com.ibm.servlet.engine.webapp.SimpleFileServlet`, to their Web module.

The Version 4.x single path limitation does not exist in Version 5. However, users who choose to use the `com.ibm.servlet.engine.webapp.SimpleFileServlet` in Version 5 must do one of the following:

- Rename `com.ibm.servlet.engine.webapp.SimpleFileServlet` to `com.ibm.ws.webcontainer.servlet.SimpleFileServlet`.
- Open the EAR file in the Application Assembly Tool (AAT) and enable the `SimpleFileServlet` static file setting.

The following list identifies the other internal servlets affected by the Version 5 package name change:

- `DefaultErrorReporter`
- `AutoInvoker`

Use the Version 5 package name, `com.ibm.ws.webcontainer.servlet.<i>servlet class name</i>` for these servlets.

2. Use the WASPostUpgrade tool to migrate servlets that extend `PageListServlet` and rely on configuration information in the associated XML servlet configuration file.

In Version 4.x, the XML servlet configuration file provides configuration data for page lists and augments servlet configuration information. This file is named as either `<i>servlet_class_name</i>.servlet` or `<i>servlet_name</i>.servlet`, and is stored in the same directory as the servlet class file.

The XML servlet configuration file is not supported in WebSphere Application Server Version 5.

3. Set a content type if your servlet generates Hyper Text Markup Language (HTML) output.

The default behavior of the Web container changed in WebSphere Application Server Version 5. If the servlet developer does not specify a content type in the servlet then the container is forbidden to set one automatically. Without an explicit content type setting, the content type is set to null. The Netscape browser displays HTML source as plain text with a null content type setting.

To resolve this problem, do one of the following:

- Explicitly set a content type in your servlet.
- Open the WAR file in the Application Assembly Tool (AAT) and enable the `autoResponseEncoding` static file setting.

4. Set the Java environment variable, `com.ibm.websphere.sendredirect.compatibility`, to **true** if you want your URLs interpreted relative to the application root.

The default value of the Java environment variable `com.ibm.websphere.sendredirect.compatibility` changed in WebSphere Application Server Version 5. In Version 4, the default setting of this variable is true. In Version 5, the setting is false.

When this variable is set to false, if a URL has a leading slash, the URL is interpreted relative to the Web module/application root. However, if the URL does not have a leading slash, it is interpreted relative to the Web container root (also known as the Web server document root). Therefore, if an application has a WAR file that has a context root of myPledge\_app and a servlet that has a servlet mapping of /Intranet/, a JSP file in the WAR file cannot access the servlet when its `encodeRedirectURL` is set to /Intranet/myPledge. The JSP file can access the servlet if the `encodeRedirectURL` is set to `myPledge_app/Intranet/myPledge`, or if the `com.ibm.websphere.sendredirect.compatibility` variable is set to true.

See the [Setting the sendredirect variable](#) article for more information.

5. Use the WASPostUpgrade tool to migrate WebSphere Version 4.x enterprise applications to Version 5.

**Note:** The WebSphere Application Server Version 4.x JSP page implementation class files are not compatible with the WebSphere Application Server Version 5 JSP container.

The WASPostUpgrade tool automatically precompiles JSP files, which ensures the JSP page implementation class files are compatible with Version 5.

If you install Version 4.x EAR files, deployed with the JSP Precompile option, in Version 5, and you choose not to follow the migration path, do one of the following:

- Select the Pre-compile JSP option in the administrative console Install New Application window.  
See article [Installing a new application](#) for more information.
- Specify the `-preCompileJSPs` option when using the Wsadmin tool.

6. Import your classes if your application uses unnamed packages.

Section 8.2 of the JSP 1.2 specification states:

The JSP container creates a JSP page implementation class for each JSP page. The name of the JSP page implementation class is implementation dependent. The JSP page implementation object belongs to an implementation-dependent named package. The package used may vary between one JSP and another, so minimal assumptions should be made. The unnamed package should not be used without an explicit `import` of the class.

For example, if `myBeanClass` is in the unnamed package, and you reference it in a `jsp:useBean` tag, then you must explicitly import `myBeanClass` with the page directive `import` attribute, as shown in the following example:

```
<%@page import="myBeanClass" %>
.
.
.
<jsp:useBean id="myBean" class="myBeanClass" scope="session"/>
```

In WebSphere Application Server Version 5, the JSP engine creates JSP page implementation classes in the `org.apache.jsp` package. If a class in the unnamed package is not explicitly imported, then the javac compiler assumes the class is in package `org.apache.jsp`, and the compilation fails.

**Note:** Avoid using the unnamed package altogether because of a change made in JDK 1.4 that will affect the JSP 2.0 specification. WebSphere Application Server Version 5 ships with JDK 1.3.1, so this is not an issue with the Version 5 JSP engine, but it will become an issue in future releases.

The *Incompatibilities* section of the version 1.4 Java 2 Platform, Standard Edition (J2SE) documentation states:

The compiler now rejects import statements that import a type from the unnamed namespace. Previous versions of the compiler would accept such import declarations, even though they were arguably not allowed by the language (because the type name appearing in the import clause is not in scope). The specification is being clarified to state clearly that you cannot have a simple name in an import statement, nor can you import from the unnamed namespace.

To summarize, the syntax:

```
import SimpleName;
```

is no longer legal. Nor is the syntax

```
import ClassInUnnamedNamespace.Nested;
```

which would import a nested class from the unnamed namespace.

To fix such problems in your code, move all of the classes from the unnamed namespace into a named namespace.

### What to do next

See "Resources for learning" for links to the J2SE, JSP, and Servlet specification documentation.

---

## Default Application

The IBM WebSphere Application Server provides a default configuration that allows administrators to easily verify that the Application Server is running. When the product is installed, it includes an application server called *server1* and an enterprise application called *Default Application*.

*Default Application* contains a Web Module called *DefaultWebApplication* and an enterprise bean JAR file called *Increment*. The *Default Application* provides a number of servlets, described below. These servlets are available in the product.

For additional code examples, visit the Samples Gallery. Learn how to locate and install the Samples Gallery by viewing the Samples Gallery reference page.

The URL for accessing Samples is: <http://localhost:9080/WSamples/>

## Snoop

Use the Snoop servlet to retrieve information about a servlet request. This servlet returns the following information:

- Servlet initialization parameters
- Servlet context initialization parameters
- URL invocation request parameters
- Preferred client locale
- Context path
- User principal
- Request headers and their values
- Request parameter names and their values
- HTTPS protocol information
- Servlet request attributes and their values
- HTTP session information
- Session attributes and their values

The Snoop servlet includes security configuration so that when WebSphere Security is enabled, clients must supply a user ID and password to execute the servlet.

The URL for the Snoop servlet is: `http://localhost:9080/snoop/`.

## HelloHTML

Use the HelloHTML pervasive servlet to exercise the PageList support provided by the WebSphere Web container. This servlet extends the PageListServlet, which provides APIs that allow servlets to call other Web resources by name or, when using the *Client Type detection* support, by type.

You can invoke the Hello servlet from an HTML browser, speech client, or most Wireless Application Protocol (WAP) enabled browsers using the URL: `http://localhost:9080/HelloHTML.jsp`.

## HitCount

Use the HitCount Demonstration application to demonstrate incrementing a counter using a variety of methods, including:

- A servlet instance variable
- An HTTP session
- An enterprise bean

You can instruct the servlet to execute any of these methods within a transaction that you can commit or roll back. If the transaction is committed, the counter is incremented. If the transaction is rolled back, the counter is not incremented.

The enterprise bean method uses a Container-Managed Persistence enterprise bean that persists the counter value to a Cloudscape database. This enterprise bean is configured to use the Default Datasource, which is set to the DefaultDB database.

When using the enterprise bean method, you can instruct the servlet to look up the enterprise bean, either in the WebSphere global namespace, or in the namespace local to the application.

The URL for the HitCount application is: `http://localhost:9080/HitCount.jsp`.

---

## Servlets

Servlets are Java programs that use the Java Servlet Application Programming Interface (API). You must package servlets in a Web ARchive (WAR) file or Web module for deployment to the application server.

*Servlets* run on a Java-enabled Web server and extend the capabilities of a Web server, similar to the way applets run on a browser and extend the capabilities of a browser.

Servlets can support dynamic Web page content, provide database access, serve multiple clients at one time, and filter data.

For the purposes of IBM WebSphere Application Server, discussions of servlets focus on Hyper Text Transfer Protocol (HTTP) servlets, which serve Web-based clients.

---

## Developing servlets with WebSphere Application Server extensions

Several WebSphere Application Server extensions are provided for enhancing your servlets. This task provides a summary of the extensions that you can utilize.

### Steps for this task

1. Review the supported specifications.

Create Java components, referring to the Servlet specifications from Sun Microsystems.

See Resources for learning for links to coding specifications and examples.

The application server includes its own packages that extend and add to the Java Servlet Application Programming Interface (API). These extensions and additions make it easier to manage session states, create personalized Web pages, generate better servlet error reports, and access databases. Locate the Javadoc for the application server APIs in the product `install_root\web\apidocs` directory.

All the public WebSphere Application Server APIs are located in the `com.ibm.websphere...` packages.

2. Use your favorite integrated development environment (IDE), or a text editor, to develop or migrate code artifacts that meet the specifications.
3. Test the code artifacts.

### What to do next

Assemble your code artifacts into a Web module as a prerequisite to deploying the code to the application server.

## Application lifecycle listeners and events

Application lifecycle listeners and events, now part of the Servlet API, enable you to notify interested listeners when servlet contexts and sessions change. For example, you can notify users when attributes change and if sessions or servlet contexts are created or destroyed.

The lifecycle listeners give the application developer greater control over interactions with `ServletContext` and `HttpSession` objects. Servlet context listeners manage resources at an application level. Session listeners manage resources associated with a series of requests from a single client. Listeners are available for lifecycle events and for attribute modification events. The listener developer creates a class that implements the `javax` listener interface, corresponding to the desired listener functionality.

At application startup time, the container uses introspection to create an instance of your listener class and registers it with the appropriate event generator.

When a servlet context is created, the `contextInitialized` method of your listener class is invoked, which creates the database connection for the servlets in your application to use, if this context is for your application.

When the servlet context is destroyed, your `contextDestroyed` method is invoked, which releases the database connection, if this context is for your application.



## Listener classes for servlet context and session changes

The following methods are defined as part of the `javax.servlet.ServletContextListener` interface:

- `void contextInitialized(ServletContextEvent)` - Notification that the Web application is ready to process requests.

Place code in this method to see if the created context is for your Web application and if it is, allocate a database connection and store the connection in the servlet context.

- `void contextDestroyed(ServletContextEvent)` -Notification that the servlet context is about to shut down.

Place code in this method to see if the created context is for your Web application and if it is, close the database connection stored in the servlet context.

Two new listener interfaces are defined as part of the `javax.servlet` package:

- `ServletContextListener`
- `ServletContextAttributeListener`

One new filter interface is defined as part of the `javax.servlet` package:

- `FilterChain` interface - methods: `doFilter()`

Two new event classes are defined as part of the `javax.servlet` package:

- `ServletContextEvent`
- `ServletContextAttributeEvent`

Three new listener interfaces are defined as part of the `javax.servlet.http` package:

- `HttpSessionListener`
- `HttpSessionAttributeListener`
- `HttpSessionActivationListener`

One new event class is defined as part of the `javax.servlet.http` package:

- `HttpSessionEvent`

### Example: `com.ibm.websphere.DBConnectionListener.java`

The following example shows how to create a servlet context listener:

```
package com.ibm.websphere;

import java.io.*;
import javax.servlet.*;

public class DBConnectionListener implements ServletContextListener
{
    // implement the required context init method
    void contextInitialized(ServletContextEvent sce)
    {
    }

    // implement the required context destroy method
    void contextDestroyed(ServletContextEvent sce)
    {
    }
}
```

## Servlet filtering

Servlet filtering is an integral part of the Servlet 2.3 API. Servlet filtering provides a new type of object called a *filter* that can transform a request or modify a response.

You can chain filters together so that a group of filters can act on the input and output of a specified resource or group of resources.

Filters typically include logging filters, image conversion filters, encryption filters, and Multipurpose Internet Mail Extensions (MIME) type filters (functionally equivalent to the servlet chaining). Although filters are not servlets, their lifecycle is very similar.

Filters are handled in the following manner:

- The Web container determines whether it needs to construct a `FilterChain` containing the `LoggingFilter` for the requested resource.  
The `FilterChain` begins with the invocation of the `LoggingFilter` and ends with the invocation of the requested resource.
- If other filters need to go in the chain, the Web container places them after the `LoggingFilter` and before the requested resource.
- The Web container then instantiates and initializes the `LoggingFilter` (if it was not done previously) and invokes its `doFilter(FilterConfig)` method to start the chain.
- The `LoggingFilter` preprocesses the request and response objects and then invokes the filter chain `doFilter(ServletRequest, ServletResponse)` method. This method passes the processing to the next resource in the chain (in this case, the requested resource).
- Upon return from the filter chain `doFilter(ServletRequest, ServletResponse)` method, the `LoggingFilter` performs post-processing on the request and response object before sending the response back to the client.

## Filter, FilterChain, FilterConfig classes for servlet filtering

The following interfaces are defined as part of the `javax.servlet` package:

- `Filter` interface - methods: `doFilter()`, `getFilterConfig()`, `setFilterConfig()`
- `FilterChain` interface - methods: `doFilter()`
- `FilterConfig` interface - methods: `getFilterName()`, `getInitParameter()`, `getInitParameterNames()`, `getServletContext()`

The following classes are defined as part of the `javax.servlet.http` package:

- `HttpServletRequestWrapper` - methods: See the Servlet 2.3 Specification
- `HttpServletResponseWrapper` - methods: See the Servlet 2.3 Specification

## Example: `com.ibm.websphere.LoggingFilter.java`

The following example shows how to implement a filter:

```
package com.ibm.websphere;

import java.io.*;
import javax.servlet.*;

public class LoggingFilter implements Filter
{
    File _loggingFile = null;
```

```

// implement the required init method
public void init(FilterConfig fc)
{
    // create the logging file
    xxx;
}

// implement the required doFilter method...this is where most of the
work is done public void doFilter(ServletRequest request, ServletResponse
response, FilterChain chain)
{
    try
    {
        // add request info to the log file
        synchronized(_loggingFile)
        {
            xxx;
        }

        // pass the request on to the next resource in the chain
        chain.doFilter(request, response);
    }
    catch (Throwable t)
    {
        // handle problem...
    }
}

// implement the required destroy method
public void destroy()
{
    // make sure logging file is closed
    _loggingFile.close();
}
}

```

## Configuring page list servlet client configurations

You can define PageListServlet configuration information in the IBM Web Extensions file. The IBM Web Extensions file is created and stored in the Web Applications archive (WAR) file by the IBM WebSphere Application Assembly Tool (AAT)

To configure and implement page lists:

### Steps for this task

1. Use the *PageList Extensions* tab in the Application Assembly Tool (AAT) to configure page list information.
2. Add the `callPage()` method to your servlet to invoke a JavaServer Page (JSP) file in response to a client request.

The PageListServlet has a `callPage()` method that invokes a JSP file in response to the HTTP request for a page in a page list. The `callPage()` method can be invoked in one of the following ways:

- `callPage(String pageName, HttpServletRequest request, HttpServletResponse response)`

where the method arguments are:

- `pageName` - a page name defined in the PageListServlet configuration
- `request` - the `HttpServletRequest` object
- `response` - the `HttpServletResponse` object

- `callPage(String mlName, String pageName, HttpServletRequest request, HttpServletResponse response)`

where the method arguments are:

- `mlName` - a markup language type
  - `pageName` - a page name defined in the `PageListServlet` configuration
  - `request` - the `HttpServletRequest` object
  - `response` - the `HttpServletResponse` object
3. Use the `PageListServlet` client type detection support to determine the markup language type a calling client requires for the response.

## Page lists

Page lists allow you to avoid hardcoding URLs in servlets and JSP files. A page list specifies the location where a request is to be forwarded, but automatically tailors that location depending on the MIME type of the servlet. These properties allow you to specify a markup language and an associated MIME type. For the given MIME type, you also specify a set of pages to invoke.

WebSphere Application Server supplies the `PageListServlet`, which you can use to call a JavaServer Pages (JSP) file by name based on the configuration data in the `client_types.xml` file. This file maps a JSP file to a Uniform Resource Identifier (URI). When the URI is invoked, it specifies another JSP file in a Web module. This support allows you to access multiple Uniform Resource Locators (URLs) without hard-coding them in your servlets.

You can also logically group page lists according to the markup language type, as for example, Hypertext Markup Language (HTML) or Wireless Markup Language (WML). This allows applications, using servlets that extend the `PageListServlet`, to call JSP files that return the proper markup-language type for the client request. For example, if a request originates from a PDA device that requires WML data and is sent to a servlet that extends the `PageListServlet`, the servlet can call a JSP file that returns a WML response.

## Client type detection support

In addition to providing the page list mapping capability, the `PageListServlet` also provides *Client Type Detection* support. A servlet determines the markup language type that a calling client needs in the response, using the configuration information in the `client_types.xml` file.

Client type detection support allows a servlet, extending the `PageListServlet`, to call an appropriate JavaServer Pages (JSP) file. The servlet invokes the `callPage()` method, which calls a JSP file based on the markup-language type of the request.

### `client_types.xml`

The `client_types.xml` file provides client type detection support for servlets extending `PageListServlet`. Using the configuration data in the `client_types.xml` file, servlets can determine the language type that calling clients require for the response.

The client type detection support allows servlets to call appropriate JavaServer Pages (JSP) files with the `callPage()` method. Servlets select JSP files based on the markup-language type of the request.

Servlets must use the following version of the `callPage()` method to determine the markup language type required by the client:

```
callPage(String mlName, String pageName, HttpServletRequest request, HttpServletResponse response)
```

where the arguments are:

- mlName - a markup language type
- pageName - a page name defined in the PageListServlet configuration
- request - the HttpServletRequest object
- response - the HttpServletResponse object

Review the Extending PageListServlet code example to see how the callPage() method is invoked by a servlet.

In the example, the client type detection method, getMLTypeFromRequest(HttpServletRequest request), provided by the PageListServlet, inspects the HttpServletRequest object request headers, and searches for a match in the client\_types.xml file.

The client type detection method does the following:

- Uses the input HttpServletRequest and the client\_types.xml file, to check for a matching HTTP request name and value.
- Returns the markup-language value configured for the <client-type> element, if a match is found.  
If multiple matches are found, this method returns the markup-language for the first <client-type> element for which a match is found.
- If no match is found, returns the value of the markup-language for the default page defined in the PageListServlet configuration.

### Location

The client\_types.xml file is located in the <i>install\_root</i>/properties directory.

### Usage notes

- Is this file read-only?  
No
- Is this file updated by a product component?  
No
- If so, what triggers its update?  
This file is created and updated manually by users.
- How and when are the contents of this file used?  
Servlets, extending PageListServlet, use this file to determine the language type that calling clients require for the response.

### Sample file entry

```
<?xml version="1.0" >
<!DOCTYPE clients [
<!ELEMENT client-type (description, markup-language, request-header+)>
<!ELEMENT description (#PCDATA)>
<!ELEMENT markup-language (#PCDATA)>
<!ELEMENT request-header (name, value)>
<!ELEMENT clients (client-type+)>
<!ELEMENT name (#PCDATA)>
<!ELEMENT value (#PCDATA)>]>
<clients>
  <client-type>
    <description>IBM Speech Client</description>
    <markup-language>VXML</markup-language>
```

```

<request-header>
  <name>user-agent</name>
  <value>IBM VoiceXML pre-release version 000303</value>
</request-header>
<request-header>
  <name>accept</name>
  <value>text/vxml</value>
</request-header>
</client-type>
<client-type>
  <description>WML Browser</description>
  <markup-language>WML</markup-language>
<request-header>
  <name>accept</name>
  <value>text/x-wap.wml</value>
</request-header>
<request-header>
  <name>accept</name>
  <value>text/vnd.wap.xml</value>
</request-header>
</client-type>
</clients>

```

### Example: Extending PageListServlet

The following example shows how a servlet extends the PageListServlet class and determines the markup-language type required by the client. The servlet then uses the callPage() method to call an appropriate JavaServer Pages (JSP) file. In this example, the JSP file that provides the the correct markup-language for the response is *Hello.page*.

```

public class HelloPervasiveServlet extends PageListServlet implements Serializable
{
  /*
  * doGet -- Process incoming HTTP GET requests
  */
  public void doGet(HttpServletRequest request, HttpServletResponse response)
  throws IOException, ServletException
  {
    // This is the name of the page to be called:
    String pageName = "Hello.page";

    // First check if the servlet was invoked with a queryString that contains a
    // markup-language value.
    // For example, if this is how the servlet is invoked:
    // http://localhost/servlets/HelloPervasive?mlname=VXML
    // then use the following method:
    String mlname= getMLNameFromRequest(request);

    // If no markup language type is provided in the queryString, then try to
    // determine the client type from the request, and use the markup-language
    // name configured in the client_types.xml file.
    if (mlName == null)
    {
      mlName = getMLTypeFromRequest(request);
    }
    try
    {
      // Serve the request page.
      callPage(mlName, pageName, request, response);
    }
    catch (Exception e)
    {
      handleError(mlName, request, response, e);
    }
  }
}

```

## autoRequestEncoding and autoResponseEncoding

Two new WebSphere Application Server extensions are available in Version 5, `autoRequestEncoding` and `autoResponseEncoding`.

In WebSphere Application Server Version 5, the Web container no longer automatically sets request and response encodings, and response content types. Programmers are expected to set these values using available methods in the Servlet 2.3 Specification. If programmers choose not to use the character encoding methods, they can specify the `autoRequestEncoding` and `autoResponseEncoding` extensions, which enable the application server to set the encoding values and content type.

The values of the `autoRequestEncoding` and `autoResponseEncoding` extensions are either `true` or `false`. The default value for both extensions is `false`. If the value is `false` for both `autoRequestEncoding` and `autoResponseEncoding`, then the request and response character encoding is set to the Servlet 2.3 Specification default, which is ISO-8859-1. Also, If the value is set to `false` for a response, the Web container cannot set a response content type.

Use the Application Assembly Tool (AAT) to change the default values for the `autoRequestEncoding` and `autoResponseEncoding` extensions.

Review the `autoRequestEncoding` and `autoResponseEncoding` encoding examples for a description of Web container behavior when these values are set to `true`.

## autoRequestEncoding and autoResponseEncoding encoding examples

The default value of the `autoRequestEncoding` and `autoResponseEncoding` extensions is `false`, which means that both the request and response character encoding is set to the Servlet 2.3 Specification default of ISO-8859-1. Different character encodings are possible if the client defines character encoding in the request header, or if the code includes the `setCharacterEncoding(String encoding)` method. Also, If the value is set to `false` for a response, the Web container cannot set a response content type.

If the `autoRequestEncoding` value is set to `true`, and the client did not specify character encoding in the request header, and the code does not include the `setCharacterEncoding(String encoding)` method, the Web container tries to determine the correct character encoding for the request parameters and data.

The Web container performs each step in the following list until a match is found:

- Looks at the character set (charset) in the *Content-Type* header.
- Attempts to map the servers locale to a character set using defined properties.
- Attempts to use the `DEFAULT_CLIENT_ENCODING` system property, if one is set.
- Uses the ISO-8859-1 character encoding as the default.

If the `autoResponseEncoding` value is set to `true`, and the client did not specify character encoding in the request header, and the code does not include the `setCharacterEncoding(String encoding)` method, the Web container does the following:

- Attempts to determine the response content type and character encoding from information in the request header.
- Uses the ISO-8859-1 character encoding as the default.

---

## JavaServer Pages files

JavaServer Pages (JSP) files are application building blocks coded to the Sun Microsystems JavaServer Pages (JSP) Specification. JSP files enable the separation of the Hypertext Markup Language (HTML) code from the business logic in Web pages so that HTML programmers and Java programmers can more easily collaborate in creating and maintaining pages.

The IBM extensions to the JSP Specification include JSP tags that resemble HTML tags making it easy for HTML authors to add the power of Java technology to Web pages, without being experts in Java programming.

JSP files support a division of roles:

### **HTML authors**

Develop JSP files that access databases and reusable Java components, such as servlets and beans.

### **Java programmers**

Create the reusable Java components and provide the HTML authors with the component names and attributes.

### **Database administrators**

Provide the HTML authors with the name of the database access and table information.

---

## Developing JavaServer Pages files with WebSphere extensions

Several IBM WebSphere extensions are provided for enhancing your JavaServer Pages (JSP) files. This task provides a summary of the extensions that you can utilize.

### Steps for this task

1. Review the supported specifications.

Create Java components, referring to the JSP specifications from Sun Microsystems.

See Resources for learning for links to coding specifications and examples.

WebSphere Application Server Version 3.5 added IBM extensions to the base Application Programming Interfaces (APIs). Since the JavaServer Pages (JSP) 1.1 and JSP 1.2 Specifications are backward compatible to the JSP 1.0 Specifications, you can invoke the APIs with the IBM extensions without modification.

The extensions belong to these categories:

### **Syntax for variable data**

Put variable fields in JSP files and have servlets and beans dynamically replace the variables with values from a database when the JSP output is returned to the browser.

### **Syntax for database access**

Add a database connection to a Web page and then use that connection to query or update the database. You can provide the user ID and password for the database connection at request time, or you can hard code the user ID and password within the JSP file.

2. Use your favorite integrated development environment (IDE), or a text editor, to develop or migrate code artifacts that meet the specifications.
3. Test the code artifacts.



4. (Optional) Batch compile your JSP files if necessary.

## Tag libraries

Java ServerPages (JSP) tag libraries contain classes for common tasks such as processing forms and accessing databases from JSP files.

Tag libraries encapsulate, as simple tags, core functionality common to many Web applications. The Java Standard Tag Library (JSTL) supports common programming tasks such as iteration and conditional processing, and provides tags for:

- manipulating XML documents
- supporting internationalization
- using Structured Query Language (SQL)

Tag libraries also introduce the concept of an expression language to simplify page development, and include a version of the JSP expression language.

A tag library has two parts - a Tag Library Descriptor (TLD) file and a JAR file.

## tsx:dbconnect tag JavaServer Pages syntax

Use the `<tsx:dbconnect>` tag to specify information needed to make a connection to a Java Database Connectivity (JDBC) or an Open Database Connectivity (ODBC) database.

The `<tsx:dbconnect>` syntax does not establish the connection. Use the `<tsx:dbquery>` and `<tsx:dbmodify>` syntax instead to reference a `<tsx:dbconnect>` tag in the same JavaServer Pages (JSP) file to establish the connection.

When the JSP file compiles into a servlet, the Java processor adds the Java coding for the `<tsx:dbconnect>` syntax to the servlet service() method, which means a new database connection is created for each request for the JSP file.

This section describes the syntax of the `<tsx:dbconnect>` tag.

```
<tsx:dbconnect id="connection_id"
  userid="db_user" passwd="user_password"
  url="jdbc:subprotocol:database"
  driver="database_driver_name"
  jndiname="JNDI_context/logical_name">
</tsx:dbconnect>
```

where:

- **id**

Represents a required identifier. The scope is the JSP file. This identifier is referenced by the connection attribute of a `<tsx:dbquery>` tag.

- **userid**

Represents an optional attribute that specifies a valid user ID for the database that you want to access. Specify this attribute to add the attribute and its value to the request object.

Although the `userid` attribute is optional, you must provide the user ID. See `<tsx:userid>` and `<tsx:passwd>` for an alternative to hard coding this information in the JSP file.

- **passwd**

Represents an optional attribute that specifies the user password for the `userid` attribute. (This attribute is not optional if the `userid` attribute is specified.) If you specify this attribute, the attribute and its value are added to the request object. Although the `passwd` attribute is optional, you must provide the password. See `<tsx:userid>` and `<tsx:passwd>` for an alternative to hard coding this attribute in the JSP file.

- **url and driver**

Represents a required attribute if you want to establish a database connection. You must provide the URL and driver.

The application server supports connection to JDBC databases and ODBC databases.

- For a JDBC database, the URL consists of the following colon-separated elements: `jdbc`, the subprotocol name, and the name of the database to access. An example for a connection to the Sample database included with IBM DB2 is:

```
url="jdbc:db2:sample"  
driver="COM.ibm.db2.jdbc.app.DB2Driver"
```

- For an ODBC database, use the Sun JDBC-to-ODBC bridge driver included in their Java2 Software Developers Kit (SDK) or another vendor's ODBC driver. The `url` attribute specifies the location of the database. The `driver` attribute specifies the name of the driver to use in establishing the database connection.

If the database is an ODBC database, you can use an ODBC driver or the Sun JDBC-to-ODBC bridge. If you want to use an ODBC driver, refer to the driver documentation for instructions on specifying the database location with the `url` attribute and the driver name.

If you use the bridge, the `url` syntax is `jdbc:odbc:<i>database</i>`. An example follows:

```
url="jdbc:odbc:autos"  
driver="sun.jdbc.odbc.JdbcOdbcDriver"
```

**Note:** To enable the application server to access the ODBC database, use the ODBC Data Source Administrator to add the ODBC data source to the System DSN configuration. To access the ODBC Administrator, click the ODBC icon on the Windows NT Control Panel.

- **jndiname**

Represents an optional attribute that identifies a valid context in the application server Java Naming and Directory Interface (JNDI) naming context and the logical name of the data source in that context. The Web administrator configures the context using an administrative client such as the WebSphere Administrative Console.

If you specify the `jndiname` attribute, the JSP processor ignores the `driver` and `url` attributes on the `<tsx:dbconnect>` tag.

An empty element (such as `<url></url>`) is valid.

## dbquery tag JavaServer Pages syntax

Use the `<tsx:dbquery>` tag to establish a connection to a database, submit database queries, and return the results set.

The `<tsx:dbquery>` tag does the following:

1. References a `<tsx:dbconnect>` tag in the same JavaServer Pages (JSP) file and uses the information the tag provides to determine the database URL and driver. You can also obtain the user ID and password from the `<tsx:dbconnect>` tag if those values are provided in the `<tsx:dbconnect>` tag.
2. Establishes a new connection
3. Retrieves and caches data in the results object.
4. Closes the connection and releases the connection resource.

This section describes the syntax of the `<tsx:dbquery>` tag.

```
<%-- SELECT commands and (optional) JSP syntax can be placed within the tsx:dbquery. --%>
<%-- Any other syntax, including HTML comments, are not valid. --%>
<tsx:dbquery id="query_id" connection="connection_id" limit="value" >
</tsx:dbquery>
```

where:

- **id**

Represents the identifier of this query. The scope is the JSP file. Use `id` to reference the query. For example, from the `<tsx:getProperty>` tag, use `id` to display the query results.

The `id` becomes the name of a bean that contains the results set. The bean properties are dynamic and the property names are the names of the columns in the results set. If you want different column names, use the SQL keyword for specifying an alias on the SELECT command. In the following example, the database table contains columns named `FNAME` and `LNAME`, but the SELECT statement uses the `AS` keyword to map those column names to `FirstName` and `LastName` in the results set:

```
Select FNAME, LNAME AS FirstName, LastName from Employee where FNAME='Jim'
```

- **connection**

Represents the identifier of a `<tsx:dbconnect>` tag in this JSP file. The `<tsx:dbconnect>` tag provides the database URL, driver name, and optionally, the user ID and password for the connection.

- **limit**

Represents an optional attribute that constrains the maximum number of records returned by a query. If this attribute is not specified, no limit is used. In such a case, the effective limit is determined by the number of records and the system caching capability.

- **SELECT command and JSP syntax**

Represents the only valid SQL command, `SELECT`. The `<tsx:dbquery>` tag must return a results set. Refer to your database documentation for information about the `SELECT` command. See other articles in this section for a description of JSP syntax for variable data and inline Java code.

## dbmodify tag JavaServer Pages syntax

The `<tsx:dbmodify>` tag establishes a connection to a database and then adds records to a database table.

The `<tsx:dbmodify>` tag does the following:

1. References a `<tsx:dbconnect>` tag in the same JavaServer Pages (JSP) file and uses the information provided by that tag to determine the database URL and driver.

**Note:** You can also obtain the user ID and password from the `<tsx:dbconnect>` tag if those values are provided in the `<tsx:dbconnect>` tag.

2. Establishes a new connection.
3. Updates a table in the database.
4. Closes the connection and releases the connection resource.

This section describes the syntax of the `<tsx:dbmodify>` tag.

```
<%-- Any valid database update commands can be placed within the DBMODIFY tag. -->
<%-- Any other syntax, including HTML comments, are not valid. -->
<tsx:dbmodify connection="connection_id">
</tsx:dbmodify>
```

where:

- **connection**  
Represents the identifier of a `<tsx:dbconnect>` tag in this JSP file. The `<tsx:dbconnect>` tag provides the database URL, driver name, and (optionally) the user ID and password for the connection.
- Database commands  
Represents valid database commands. Refer to your database documentation for details

## tsx:getProperty tag JavaServer Pages syntax and examples

The `<tsx:getProperty>` tag gets the value of a bean to display in a JavaServer Pages (JSP) file.

This IBM extension of the Sun JSP `<jsp:getProperty>` tag implements all of the `<jsp:getProperty>` function and adds the ability to introspect a database bean created using the IBM extension `<tsx:dbquery>` or `<tsx:dbmodify>`.

**Note:** You cannot assign the value from this tag to a variable. The value, generated as output from this tag, displays in the browser window.

This section describes the syntax of the `<tsx:getProperty>` tag:

```
<tsx:getProperty name="bean_name"
  property="property_name" />
```

where:

- **name**  
Represents the name of the bean declared by the `id` attribute of a `<tsx:dbquery>` syntax within the JSP file. See `<tsx:dbquery>` for an explanation. The value of this attribute is case-sensitive.
- **property**  
Represents the property of the bean to access for substitution. The value of the attribute is case-sensitive and is the locale-independent name of the property.

Tag example:

```
<tsx:getProperty name="userProfile" property="username" />
<tsx:getProperty name="request" property=request.getParameter("corporation") />
```

In most cases, the value of the property attribute is just the property name. However, to access the request bean or to access a property of a property (sub property), specify the full form of the property attribute. The full form also gives you the option to specify an index for indexed properties. You can specify the

optional index as a constant (such as 2), or an index like the one described in the `<tsx:repeat>` tag. Some examples using the full form of the property attribute follow:

```
<tsx:getProperty name="staffQuery" property=address(currentAddressIndex) />
<tsx:getProperty name="shoppingCart" property=items(4).price />
<tsx:getProperty name="fooBean" property=foo(2).bat(3).boo.far />
```

## tsx:userid and tsx:passwd tag JavaServer Pages syntax

With the `<tsx:userid>` and `<tsx:passwd>` tags, you do not have to hard code a user ID and password in the `<tsx:dbconnect>` tag.

Use the `<tsx:userid>` and `<tsx:passwd>` tags to accept user input for the values and then add that data to the request object. You can access the request object with a JavaServer Pages (JSP) file, such as the *JSPEmployee.jsp* example that requests the database connection.

You must use `<tsx:userid>` and `<tsx:passwd>` tags within a `<tsx:dbconnect>` tag.

This section describes the syntax of the `<tsx:userid>` and `<tsx:passwd>` tags.

```
<tsx:dbconnect id="connection_id"
  <font color="red"><userid></font>
  <tsx:getProperty name="request" property=request.getParameter("userid")
  /><font color="red"></userid></font>
  <font color="red"><passwd></font>
  <tsx:getProperty name="request" property=request.getParameter("passwd")
  /><font color="red"></passwd></font>
  url="protocol:database_name:database_table"
  driver="JDBC_driver_name">
</tsx:dbconnect>
```

where:

- **<tsx:getProperty>**

Represents the syntax as a mechanism for embedding variable data.

- **userid**

Represents a reference to the request parameter that contains the user ID. You must add the parameter to the request object that passes to this JSP file. You can set the attribute and its value in the request object, using an HTML form or a URL query string to pass the user-specified request parameters.

- **passwd**

Represents a reference to the request parameter that contains the password. Add the parameter to the request object that passes to this JSP file. You can set the attribute and its value in the request object, using an HTML form or a URL query string, to pass user-specified values.

## tsx:repeat tag JavaServer Pages syntax

The `<tsx:getProperty>` tag repeats a block of HTML tagging.

Use the `<tsx:repeat>` syntax to iterate over a database query results set. The `<tsx:repeat>` syntax iterates from the start value to the end value until one of the following conditions is met:

- The end value is reached.
- An exception is thrown.

The output of a `<tsx:repeat>` block is buffered until the block completes. If an exception is thrown before a block completes, no output is written for that block.

This section describes the syntax of the `<tsx:repeat>` tag:

```
<tsx:repeat index=name start="starting_index" end="ending_index">
</tsx:repeat>
```

where:

- **index**  
Represents an optional name used to identify the index of this repeat block. The value is case-sensitive and its scope is the JSP file.
- **start**  
Represents an optional starting index value for this repeat block. The default is 0.
- **end**  
Represents an optional ending index value for this repeat block. The maximum value is 2,147,483,647.  
If the value of the end attribute is less than the value of the start attribute, the end attribute is ignored.

## Example: Combining `tsx:repeat` and `tsx:getProperty` JavaServer Pages tags

The following code snippet shows you how to code these tags:

```
<tsx:repeat>
<tr>
  <td><tsx:getProperty name="empqs" property="EMPNO" />
  <tsx:getProperty name="empqs" property="FIRSTNME" />
  <tsx:getProperty name="empqs" property="WORKDEPT" />
  <tsx:getProperty name="empqs" property="EDLEVEL" />
</td>
</tr>
</tsx:repeat>
```

## Example: `tsx:dbmodify` tag syntax

In the following example, a new employee record is added to a database. The values of the fields are based on user input from this JavaServer Pages (JSP) file and referenced in the database commands using the `<tsx:getProperty>` tag.

```
<tsx:dbmodify connection="conn" >
insert into EMPLOYEE
  (EMPNO,FIRSTNME,MIDINIT,LASTNAME,WORKDEPT,EDLEVEL)
values
('<tsx:getProperty name="request" property=request.getParameter("EMPNO") />',
'<tsx:getProperty name="request" property=request.getParameter("FIRSTNME") />',
'<tsx:getProperty name="request" property=request.getParameter("MIDINIT") />',
'<tsx:getProperty name="request" property=request.getParameter("LASTNAME") />',
'<tsx:getProperty name="request" property=request.getParameter("WORKDEPT") />',
<tsx:getProperty name="request" property=request.getParameter("EDLEVEL") />)
</tsx:dbmodify>
```

## Example: Using `tsx:repeat` JavaServer Pages tag to iterate over a results set

The `<tsx:repeat>` tag iterates over a results set. The results set is contained within a bean. The bean can be a static bean, for example, a bean created by using the IBM WebSphere Studio database wizard, or a dynamically generated bean, for example,

a bean generated by the `<tsx:dbquery>` syntax. The following table is a graphic representation of the contents of a bean called, *myBean*:

|      | col1    | col2    | col3       |
|------|---------|---------|------------|
| row0 | friends | Romans  | countrymen |
| row1 | bacon   | lettuce | tomato     |
| row2 | May     | June    | July       |

Some observations about the bean:

- The column names in the database table become the property names of the bean. The `<tsx:dbquery>` section describes a technique for mapping the column names to different property names.
- The bean properties are indexed. For example, `myBean.get(Col1(row2))` returns May.
- The query results are in the rows. The `<tsx:repeat>` tag iterates over the rows, beginning at the start row.

The following table compares using the `<tsx:repeat>` tag to iterate over a static bean, versus a dynamically generated bean:

| Static Bean Example  | <code>&lt;tsx:repeat&gt;</code> Bean Example   |
|--|--|
| <p><b>myBean.class</b></p> <pre>// Code to get a connection  // Code to get the data Select * from myTable;  // Code to close the connection</pre> <p><b>JSP file</b></p> <pre>&lt;tsx:repeat index=abc&gt;   &lt;tsx:getProperty name="myBean"     property="col1(abc)" /&gt; &lt;/tsx:repeat&gt;</pre> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• The bean (<code>myBean.class</code>) is a static bean.</li> <li>• The method to access the bean properties is <code>myBean.get(property(index))</code>.</li> <li>• You can omit the property index, in which case the index of the enclosing <code>&lt;tsx:repeat&gt;</code> tag is used. You can also omit the index on the <code>&lt;tsx:repeat&gt;</code> tag.</li> <li>• The <code>&lt;tsx:repeat&gt;</code> tag iterates over the bean properties row by row, beginning with the start row.</li> </ul> | <p><b>JSP file</b></p> <pre>&lt;tsx:dbconnect id="conn" userid="alice"passwd="test" url="jdbc:db2:sample" driver="COM.ibm.db2.jdbc.app.DB2Driver"&gt; &lt;/tsx:dbconnect &gt;  &lt;tsx:dbquery id="dynamic" connection="conn" &gt;   Select * from myTable; &lt;/tsx:dbquery&gt;  &lt;tsx:repeat index=abc&gt;   &lt;tsx:getProperty name="dynamic"     property="col1(abc)" /&gt; &lt;/tsx:repeat&gt;</pre> <p><b>Notes:</b></p> <ul style="list-style-type: none"> <li>• The bean (dynamic) is generated by the <code>&lt;tsx:dbquery&gt;</code> tag and does not exist until the syntax executes.</li> <li>• The method to access the bean properties is <code>dynamic.getValue("property", index)</code>.</li> <li>• You can omit the property index, in which case the index of the enclosing <code>&lt;tsx:repeat&gt;</code> tag is used. You can also omit the index on the <code>&lt;tsx:repeat&gt;</code> tag.</li> <li>• The <code>&lt;tsx:repeat&gt;</code> tag syntax iterates over the bean properties row by row, beginning with the start row.</li> </ul> |

## Implicit and explicit indexing

Examples 1, 2, and 3 show how to use the `<tsx:repeat>` tag. The examples produce the same output if all indexed properties have 300 or fewer elements. If there are more than 300 elements, Examples 1 and 2 display all elements, while Example 3 shows only the first 300 elements.

Example 1 shows *implicit indexing* with the default start and default end index. The bean with the smallest number of indexed properties restricts the number of times the loop repeats.

```
<table>
<tsx:repeat>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="city" /></tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="address" /></tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="telephone" /></tr></td>
</tsx:repeat>
</table>
```

Example 2 shows indexing, starting index, and ending index:

```
<table>
<tsx:repeat index=myIndex start=0 end=2147483647>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property=city(myIndex) />
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property=address(myIndex) />
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property=telephone(myIndex) />
  </tr></td>
</tsx:repeat>
</table>
```

Example 3 shows *explicit indexing* and ending index with implicit starting index. Although the index attribute is specified, you can still implicitly index the indexed property `city` because the `(myIndex)` tag is not required.

```
<table>
<tsx:repeat index=myIndex end=299>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="city" /t>
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="address(myIndex)" />
  </tr></td>
  <tr><td><tsx:getProperty name="serviceLocationsQuery" property="telephone(myIndex)" />
  </tr></td>
</tsx:repeat>
</table>
```

## Nesting `<tsx:repeat>` blocks

You can nest `<tsx:repeat>` blocks. Each block is separately indexed. This capability is useful for interleaving properties on two beans, or properties that have subproperties. In the example, two `<tsx:repeat>` blocks are nested to display the list of songs on each compact disc in the user's shopping cart.

```
<tsx:repeat index=cdindex>
  <h1><tsx:getProperty name="shoppingCart" property=cds.title /></h1>
  <table>
    <tsx:repeat>
      <tr><td><tsx:getProperty name="shoppingCart" property=cds(cdindex).playlist />
      </td></tr>
    </tsx:repeat>
  </table>
</tsx:repeat>
```



## JspBatchCompiler tool

As an IBM enhancement to JavaServer Pages support, IBM WebSphere Application Server provides a batch JSP compiler. Use this function to batch compile your JSP files and thereby enable faster responses to the initial client requests for the JSP files on your production Web server.

Batch compiling makes the first request for a JSP file much faster because the JSP file is translated and compiled into a servlet. Batch compiling is also useful as a fast way to resynchronize all of the JSP files for an application.

To use the JSP batch compiler for JSP files, enter the following command on a single line at an operating system command prompt:

```
JspBatchCompiler -enterpriseapp.name <name>
                  [-webmodule.name <name>]
                  [-cell.name <name>]
                  [-node.name <name>]
                  [-server.name <name>]
                  [-filename <jsp name>]
                  [-keepgenerated <true|false>]
                  [-verbose <true|false>]
                  [-deprecation <true|false>]
```

If the names specified for these arguments are comprised of two or more words separated by spaces, you must add quotation marks around the names.

where:

- **enterpriseapp.name**  
Represents the name of the enterprise application you want to compile.
- **webmodule.name**  
Represents the name of the specific Web module that you want to compile. If this argument is not set, all Web modules in the enterprise application are compiled.
- **cell.name**  
Represents the name of the cell in which the application is deployed. The default is `BaseApplicationServerCell`.
- **node.name**  
Represents the name of the node in which the application is deployed. The default is `DefaultNode`.
- **server.name**  
Represents the name of the server in which the application is deployed. The default is `server1`.
- **filename**  
Represents the name of a single JSP file that you want to compile. If this argument is not set, all files in the Web module are compiled. Alternatively, if *filename* is set to the name of a directory, only the JSP files in that directory are compiled.
- **keepgenerated**  
Represents the option to save or erase the generated files.  
If set to yes, WebSphere Application Server saves the generated `.java` files used for compilation on your server. By default, this argument is set to no and the `.java` files are erased after the class files have compiled.
- **verbose**

Indicates the compiler should generate verbose output while compiling the generated sources.

- **deprecation**

Indicates the compiler should generate deprecation warnings while compiling the generated sources.

---

## Bean Scripting Framework

The Bean Scripting Framework (BSF) enables you to use scripting language functions in your Java server-side applications. This framework also extends scripting languages so that you can use existing Java classes and Java beans in the JavaScript language.

With BSF, you can write scripts that create, manipulate and access values from Java objects, or you can write Java programs that evaluate and access results from scripts.

WebSphere Application Server provides the Bean Scripting Framework, which consists of a BSF manager, a BSF engine, and a scripting engine.

BSF provides an access mechanism to Java objects for the scripting languages it supports, so that both the scripting language and the Java code can access code exclusive functions. The access mechanism is implemented through a registry of objects maintained by BSF.

BSF in WebSphere Application Server supports the Rhino ECMAScript.

The "Resources for Learning" article provides external BSF links that document future supported languages.

---

## Developing Web applications

### Before you begin

Design a Web application and the components that it needs.

For general Web application design information, see "Resources for learning."

There are two basic approaches to selecting tools for developing Web applications:

- You can use one of the available integrated development environments (IDEs). IDE tools automatically generate significant parts of the servlet and JavaServer Pages (JSP) code, and Hypertext Markup Language (HTML) files. They also contain integrated tools for packaging and testing the Web application components. The IBM WebSphere Application Developer product is the recommended IDE. For more information, see the documentation for that product.
- If you decide to develop Web components without an IDE, you need at least an ASCII text editor. You can also use tools available in the Java Software Development Kit (SDK) and in this product to assemble, test, and deploy the Web application components.

The following steps support the second approach, development without an IDE.

### Steps for this task

1. If necessary, migrate any pre-existing code to the required version of the servlet and JSP specification.
2. Write and compile the components of the Web application.

To access classes that were extended, compile your code using the `-classpath` option on the `javac` compiler. This option allows you to reference the `j2ee.jar` file in the product `install_root\lib` directory.

For example, to compile a servlet running on the Windows NT version of WebSphere Application Server, specify:

```
javac -classpath D:\Program Files\WebSphere\AppServer\lib\j2ee.jar MyServlet.java
```

To compile that same servlet on the Windows NT version of WebSphere Network Deployment, specify:

```
javac -classpath D:\Program Files\WebSphere\DeploymentManager\lib\j2ee.jar MyServlet.java
```

3. **(Optional)** Disable JavaServer Pages (JSP) runtime compilation, if necessary.

#### What to do next

Assemble the application components in one or more Web modules.

## Disabling JavaServer Pages run-time compilation

By default, the JavaServer Pages (JSP) engine translates a requested JSP file, compiles the `.java` file, and loads the compiled servlet into the run-time environment. In previous releases of WebSphere Application Server, if a `.class` file did not exist, the JSP engine always translated and compiled the JSP file. You had to turn off the Web applications reload capability to prevent additional translations and recompiles of the file.

With Version 5.0.1 of WebSphere Application Server, you can now change the JSP engine default behavior by indicating a JSP file should never be translated or compiled at run time, even when a `.class` file does not exist.

If run-time compilation is disabled, you must precompile the JSP files, which provides the following advantages:

- Reduces compilation related disk operations.
- Minimizes disk storage requirements necessary for handling temporary `.java` and `.class` files generated during a run-time compilation.
- Forces you to verify that a JSP file compiled successfully before deploying and installing the application in WebSphere Application Server.

You can disable run-time JSP file compilation on a global or an individual Web application basis:

- To disable the translation and compilation of JSP files for all Web applications, set the Web container Custom property `disableJspRuntimeCompilation` to `true`.

Set this property through the Web container Custom properties panel in the administrative console. To view this administrative console page, click **Servers** > **Application Servers** > *server\_name* > **Web Container** > **Custom Properties** > *property\_name*.

Valid values for this setting are `true` or `false`. If this property is set to `true`, then translation and compilation of the JSP files is disabled at run time for all Web applications.

- To disable the translation and compilation of JSP files for a specific Web application, set the JSP engine initialization parameter

`disableJspRuntimeCompilation` to true. This setting, if enabled, determines the run-time behavior of the JSP engine and overrides the Web container custom property setting.

Set this parameter through the JavaServer Pages attribute assembly settings panel in the Application Assembly Tool (AAT). To view this page, click **Web Modules** > *component\_instance* > **Assembly Property Extensions**.

Valid values for this setting are true or false. If this parameter is set to true, then, for that specific Web application, translation and compilation of the JSP files is disabled at run time, and the JSP engine only loads precompiled files.

- If neither the Web container custom property nor the JSP attribute assembly parameter is set, the first request for a JSP file results in the translation and compilation of the JSP file when the `.class` file does not exist. Subsequent requests for the file also result in compilations and translations, but only if the following conditions are met:
  - Compilations and translations are required.
  - Reloading is enabled for the Web module.
  - Reload interval is exceeded.

If you disable run-time compilation and a request arrives for a JSP file that does not have a matching `.class` file, the JSP engine returns HTTP error 501 (Not implemented) to the browser. If the JSP file does not exist, the JSP engine returns HTTP error 404 (File not found) to the browser. In both cases, an exception is written to the System Out (SYSOUT) and First Failure Data Capture (FFDC) logs. If a JSP file has a matching `.class` file but that file is out of date, the JSP engine still loads the `.class` file into memory.

Perform the following steps to determine whether the `disableJspRuntimeCompilation` option is enabled in WebSphere Application Server:

1. Enable the Diagnostic Trace Service and set the trace specification to `com.ibm.ws.webcontainer.jsp.servlet.*=all=enabled`.
2. Request a JSP file.
3. Locate the string, `disableJspRuntimeCompilation:true`, in the `trace.log` file.
4. Ensure the `jspUri`: entry matches the requested JSP file.

If both the `disableJspRuntimeCompilation:true` string and the matching `jspUri`: entry appear in the trace, the `disableJspRuntimeCompilation` setting is enabled for the Web application.

---

## Example: Converting JavaScript source to the Bean Scripting Framework

JavaScript code is one of the most popular languages of Web developers. This language supports the following base objects, plus additional objects from the Document Object Model:

- array
- date
- math
- number
- string

Server-side JavaScript code supports the same base objects, and additional objects that support user access to databases, file systems and e-mail systems.

Like client-side JavaScript code, server-side JavaScript code is also platform, browser, and language independent.

You can convert server-side JavaScript applications to the Bean Scripting Framework. This article describes how to perform this conversion.

### Server-side JavaScript source code

Suppose you have the following server-side JavaScript application:

```
<html>
<head>
<title>Hello World server-side JavaScript example</title>
</head>
<body>
<br><br>
</body>
</html>

<server>
function writePage()
  write("<center><font size='6'>Hello World</font></center>");
</server>
```

### Converting server-side JavaScript source code to the Bean Scripting Framework (BSF)

Make the following changes to the JavaScript source code to enable BSF:

```
<%@ page language="javascript" %>
<html>
<head>
<title>Hello World server-side BSF/JavaScript example</title>
</head>
<body>
<br><br>
</body>
</html>

<%
  out.println("<center><font size='6'>Hello World</font></center>");
%>
```

Review the other BSF reference articles for deployment information and additional programming examples.

---

## Scenario: Creating a Bean Scripting Framework application

### Scenario description

Programming skills in JavaScript code are more prevalent than programming skills using JavaServer Pages (JSP) tags. Using the Bean Scripting Framework, JavaScript programmers can gradually introduce JSP tags in their JavaScript applications without completely rewriting the source code. The BSF method not only reduces the potential of programming errors, but also provides a painless way to learn a new technology.

The following scenario illustrates how to implement a BSF application using JavaScript within JSP tags.

## Developing the BSF application

At ABC elementary school, John Doe teaches third grade mathematics. He wants to help his students memorize their multiplication tables, and thinks a small Web-based quiz could help meet his objective. However, John Doe only knows JavaScript.

Using the Bean Scripting Framework to help leverage his JavaScript skills, John Doe creates two JSP files, `multiplication_test.jsp` and `multiplication_scoring.jsp`.

In the `multiplication_test.jsp` file, John Doe uses both client-side and server-side JavaScript code to generate a test of 100 random multiplication questions, displayed using a three minute timer. He then writes the `multiplication_scoring.jsp` file to read the data submitted by the `multiplication_test.jsp` file and to generate the scoring results.

John Doe creates the following two files:

### **multiplication\_test.jsp:**

```
<html>
<head>
<title>Multiplication Practice Test</title>
<script language="javascript">
var countMin=3;
var countSec=0;
function updateDisplay (min, sec) {
  var disp;
  if (min <= 9) disp = " 0";
  else disp = " ";
  disp += (min + ":");
  if (sec <= 9) disp += ("0" + sec);
  else disp += sec;
  return(disp);
}
function countDown() {
  countSec--;
  if (countSec == -1) {
    countSec = 59;
    countMin--;
  }
  document.multtest.counter.value = updateDisplay(countMin, countSec);
  if((countMin == 0) &&(countSec == 0)) document.multtest.submit();
  else var down = setTimeout("countDown();", 1000);
}
</script>
</head>
<body bgcolor="#ffffff" onLoad="countDown();">
<%@ page language="javascript" %>
<h1>Three Minute Multiplication Drill</h1>
<hr>
<h2>Remember: this is an opportunity to excel!</h2>
<p>
<form method="POST" name="multtest" action="multiplication_scoring.jsp">
<div align="center">
<table>
<tr>
<td>
<h3>Time left:
<input type="text" name="counter" size="9" value="03:00" readonly>
</h3>
</td>
<td>
<input type="submit" value="Submit for scoring!">
</td>
</tr>
</table>
</div>
</body>
```

```

</tr>
</table>
<table border="1">
<%
var newrow = 0;
var q_num = 0;
function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");
    out.println("<td>");
    out.println(num1 + " x " + num2 + " = ");
    out.println("</td><td>");
    out.print("<input name=\"" + q_num + "|" + num1 + ":" + num2 + "\" ");
    out.println("type=\"text\" size=\"10\">");
    out.println("</td>");
    if (newrow == 3) {
        out.println("</tr>");
        newrow = 0;
    }
    else newrow++;
    q_num++;
}
for (var i = 0; i < 100; i++) {
    var rand1 = Math.ceil(Math.random() * 12);
    var rand2 = Math.ceil(Math.random() * 12);
    addQuestion(rand1, rand2);
}
%>
</table>
</div>
</form>
</body>
</html>

```

#### **multiplication\_scoring.jsp:**

```

<html>
<head>
<title>Multiplication Practice Test Results</title>
</head>
<body bgcolor="#ffffff">
<%@ page language="javascript" %>
<h1>Multiplication Drill Score</h1>
<hr>
<div align="center">
<table border="1">
<tr><th>Problem</th><th>Correct Answer</th><th>Your Answer</th></tr>
<%
var total_score = 0;
function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());
    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;
    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");
    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
    out.println(your_product + "</td>");
    out.println("</tr>");
}
var equations = request.getParameterNames();
while(equations.hasMoreElements()) {
    var currElt = equations.nextElement();

```

```

        var splitPos1 = currElt.indexOf("|");
        var splitPos2 = currElt.indexOf(":");
        if (splitPos1 >=0 && splitPos2 >= 0) score(currElt, splitPos1, splitPos2);
    }
    %>
</table>
<h2>Total Score: <%= total_score %></h2>
<h3><a href="multiplication_test.jsp">Try again?</a></h3>
</div>
</body>
</html>

```

Follow these steps to see how John Doe uses BSF to implement JavaScript in a JSP application:

### Steps for this task

1. Give your files a .jsp extension.
2. Use server-side JavaScript code in your application.

The `multiplication_test.jsp` file incorporates both client-side and server-side JavaScript. Server-side JavaScript is similar to client-side JavaScript; the primary difference consists of using a different set of objects. Whereas client-side Javascript programmers invoke `document` and `window` objects, server-side JavaScript programmers, using the Bean Scripting Framework, invoke a set of objects provided by the JSP technology. Also, client-side scripts are enclosed in `<script>` tags, but server-side scripts use JSP scriptlet and expression tags.

Examine the following blocks of code:

```

<script language="javascript">
var countMin=3;
var countSec=0;
function updateDisplay (min, sec) {
    var disp;
    if (min <= 9) disp = " 0";
    else disp = " ";
    disp += (min + ":");
    if (sec <= 9) disp += ("0" + sec);
    else disp += sec;
    return(disp);
}
function countDown() {
    countSec--;
    if (countSec == -1) {
        countSec = 59;
        countMin--;
    }
    document.multtest.counter.value = updateDisplay(countMin, countSec);
    if((countMin == 0) && (countSec == 0)) document.multtest.submit();
    else var down = setTimeout("countDown();", 1000);
}
</script>
....
<body bgcolor="#ffffff" onLoad="countDown();">
...
<form method="POST" name="multtest" action="multiplication_scoring.jsp">
...
<input type="text" name="counter" size="9" value="03:00" readonly>
...

```

The JavaScript code contained in the `<script>` block implements a timer set within the `<input>` field named `counter`. The `onLoad` event handler in the `<body>` tag causes the browser to load and execute the code when the the page is loaded.



The `document.mlttest.submit()` statement causes the form named `mlttest` to be submitted when the timer expires.

3. Identify the code to the BSF function.

The following code example, from the `multiplication_test.jsp` file, displays the use of a JSP directive. This directive tells the WebSphere Application Server BSF function that this file is using the JavaScript language, and that the JavaScript code is enclosed by the `<% ... %>` scriptlet tags. The `out` implicit JSP object in this code example, creates the body section of a table from 100 randomly generated questions.

```
...
<%@ page language="javascript" %>
...
<%
var newrow = 0;
var q_num = 0;

function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");

    out.println("<td>");
    out.println(num1 + " x " + num2 + " = ");
    out.println("</td><td>");
    out.print("<input name=\"" + q_num + "\" |" + num1 + ":" + num2 + "\" ");
    out.println("type=\"text\" size=\"10\">");
    out.println("</td>");

    if (newrow == 3) {
        out.println("</tr>");
        newrow = 0;
    }
    else newrow++;

    q_num++;
}

for (var i = 0; i < 100; i++) {
    var rand1 = Math.ceil(Math.random() * 12);
    var rand2 = Math.ceil(Math.random() * 12);

    addQuestion(rand1, rand2);
}

%>
...

```

4. Read the results.

To score the results of the practice drill, John Doe uses the request implicit JSP object in the `multiplication_scoring.jsp` file to obtain the POST data created within the `<form>` tags in the `multiplication_test.jsp` file.

The `multiplication_scoring.jsp` file uses the POST data to build an output file containing the original question, the student's answer, and the correct answer, and then prints the text in a table format using the `out` implicit object.

The following code example from the `multiplication_scoring.jsp` file illustrates the use of the request and `out` JSP objects:

```
...
<%@ page language="javascript" %>
...
<%
var total_score = 0;
function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());

```

```

    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;
    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");
    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
    out.println(your_product + "</td>");
    out.println("</tr>");
}
var equations = request.getParameterNames();
while(equations.hasMoreElements()) {
    var currElt = equations.nextElement();
    var splitPos1 = currElt.indexOf("|");
    var splitPos2 = currElt.indexOf(":");
    if (splitPos1 >=0 && splitPos2 >= 0) score(currElt, splitPos1, splitPos2);
}
%>
...
<h2>Total Score: <%= total_score %></h2>
...

```

**Note:** Although using separate scriptlet blocks of code for different portions of a conditional expression is common in JSP files implemented in Java, it is invalid for JSP files implemented using JavaScript through the Bean Scripting Framework. The JavaScript code must be entirely contained within the scriptlet tags.

The following code example illustrates invalid usage:

```

<% if (pass == 0) %>
    <i>pass is true</i>
<% else %>
    <i>pass is not true</i>

```

## Deploying the BSF application

You assemble and deploy BSF applications in the same manner as JSP applications. Review the (Assembling applications) article for more information.

Deploy the BSF code examples in WebSphere Application Server to view this applications processing and output. Use the following quick steps to deploy the application.

**Note:** The intent of these "quick steps" is to provide you with instant application output. However, the supported method for deployment is the same as for standard JSP files.

### Steps for this task

1. Use the DefaultApplication to add your BSF files.

Copy your .jsp files to the DefaultApplication directory: <<i>app server install directory</i>/installedApps/<<i>node name</i>/DefaultApplication.ear/DefaultApplication.war

2. Start the application server.
3. Open a browser and request your BSF application

Use the following URL to request your application:  
`http://hostName:9080/<<i>jspFileName</i>>.jsp`

---

## Example: Bean Scripting Framework code example

The following code examples show how to implement JavaScript using the Bean Scripting Framework (BSF).

For a quick demonstration of the BSF function, copy these code examples into 2 separate files, and deploy them in WebSphere Application Server using the instructions in the BSF scenario article.

### Multiplication practice test

```
<html>
<head>
<title>Multiplication Practice Test</title>
<!--
This file and its companion, multiplication_score.jsp, illustrate the
use of ECMAScript within the BSF framework. The task is a simple
timed math quiz, which is 3 minutes in duration. When the quiz ends,
the score is computed and displayed. Users are then asked if they wish
to try the quiz again.
-->

<!--
This code fragment displays and updates the quiz
countdown in client side JavaScript code.
-->
<script language="javascript">
var countMin=3;
var countSec=0;

// This code computes the current countdown time.
function updateDisplay (min, sec) {
    var disp;

    if (min <= 9) disp = " 0";
    else disp = " ";

    disp += (min + ":");

    if (sec <= 9) disp += ("0" + sec);
    else disp += sec;

    return(disp);
}

// This code fragment displays the current countdown time in the user's browser
// window, and submits the results for scoring when the countdown ends.

function countDown() {
    countSec--;
    if (countSec == -1) {
        countSec = 59;
        countMin--;
    }
    document.mulpttest.counter.value = updateDisplay(countMin, countSec);
    if((countMin == 0) && (countSec == 0)) document.mulpttest.submit();
    else var down = setTimeout("countDown();", 1000);
}

</script>
</head>
<body bgcolor="#ffffff" onLoad="countDown();">
```

```

<!--
The body of the quiz runs as JavaServer Pages (JSP) code using BSF. The
code outputs the problems in table format using the POST method and invokes the
scoring module when the user chooses to end the quiz or when
the countdown ends.
-->
<%@ page language="javascript" %>

<h1>Three Minute Multiplication Drill</h1>
<hr>

<h2>Remember: this is an opportunity to excel!</h2>
<p>

<form method="POST" name="multtest" action="multiplication_scoring.jsp">
<div align="center">
<table>
<tr>
<td>
<h3>Time left:
<input type="text" name="counter" size="9" value="03:00" readonly>
</h3>
</td>
<td>
<input type="submit" value="Submit for scoring!">
</td>
</tr>
</table>
<table border="1">
<%
var newrow = 0;
var q_num = 0;

// This code generates a new random multiplication problem up to the number
// twelve, and enters it into the table of problems.

function addQuestion(num1, num2) {
    if (newrow == 0) out.println("<tr>");

    out.println("<td>");
    out.println(num1 + " x " + num2 + " = ");
    out.println("</td><td>");
    out.print("<input name=\"\" + q_num + \"|\" + num1 + \":\" + num2 + \"\ " );
    out.println("type=\"text\" size=\"10\">");
    out.println("</td>");

    if (newrow == 3) {
        out.println("</tr>");
        newrow = 0;
    }
    else newrow++;

    q_num++;
}

//This code obtains two random operands and formats 100 quiz problems.

for (var i = 0; i < 100; i++) {
    var rand1 = Math.ceil(Math.random() * 12);
    var rand2 = Math.ceil(Math.random() * 12);

    addQuestion(rand1, rand2);
}

%>

```

```

</table>
</div>
</form>

</body>
</html>

```

### Multiplication practice test results

```

<html>
<head>
<title>Multiplication Practice Test Results</title>
</head>
<body bgcolor="#ffffff">

<!--
This JSP code is invoked when the user submits a math quiz for scoring, or when the
quiz countdown expires. The JSP code tabulates the problem list, the correct answer,
the user's answer, and scores the test. It then offers the user an
opportunity to try the quiz again.
-->
<%@ page language="javascript" %>

<h1>Multiplication Drill Score</h1>
<hr>

<div align="center">
<table border="1">
<tr><th>Problem</th><th>Correct Answer</th><th>Your Answer</th></tr>
<%
var total_score = 0;

// This code parses the submitted form, extracts the a problem generated by the
// multiplication_test.jsp file, outputs it, computes the correct answer,
// and displays this information and the user answer. The code scores
// the quiz using a running sum of correct answers.

function score (current, pos1, pos2) {
    var multiplier = current.substring(pos1 + 1, pos2);
    var multiplicand = current.substring(pos2 + 1, current.length());
    var your_product = request.getParameterValues(current)[0];
    var true_product = multiplier * multiplicand;

    out.println("<tr>");
    out.println("<td>" + multiplier + " x " + multiplicand + " = </td>");
    out.println("<td>" + true_product + "</td>");

    if (your_product == true_product) {
        total_score++;
        out.print("<td bgcolor=\"\#00ff00\">");
    }
    else {
        out.print("<td bgcolor=\"\#ff0000\">");
    }
    out.println(your_product + "</td>");
    out.println("</tr>");
}

// This is the main body of the scoring application. It parses the posted quiz,
// and calls the score() function to score remaining problems.

var equations = request.getParameterNames();
while(equations.hasMoreElements()) {
    var currElt = equations.nextElement();
    var splitPos1 = currElt.indexOf("|");
    var splitPos2 = currElt.indexOf(":");

```

```

    if (splitPos1 >=0 && splitPos2 >= 0) score(currElt, splitPos1, splitPos2);
}

%>
</table>

<h2>Total Score: <%= total_score %></h2>
<h3><a href="/multiplication_test.jsp">Try again?</a></h3>
</div>

</body>
</html>

```

---

## Web modules

A Web module represents a Web application. A Web module is created by assembling servlets, JavaServer Pages (JSP) files, and static content such as HyperText Markup Language (HTML) pages into a single deployable unit. Web modules are stored in Web archive (WAR) files, which are standard Java archive files.

A Web module contains:

- One or more servlets, JSP files, and HTML files.
- A deployment descriptor, stored in an Extensible Markup Language (XML) file. The file, named `web.xml`, declares the contents of the module. It contains information about the structure and external dependencies of Web components in the module and describes how the components are used at run time.

You can create Web modules as stand-alone applications, or you can combine Web modules with other modules to create J2EE applications. You install and run a Web module in the Web container of an application server.

---

## Assembling Web Modules

### Before you begin

If you want to use existing Java 2 Enterprise Edition (J2EE) 1.2 Web modules in your J2EE 1.3 application, migrate them to J2EE 1.3 first.

Assemble a Web module to contain servlets, JSP files, and related code artifacts. (Group enterprise beans, client code, and resource adapter code in separate modules).

A Web module can be installed as a stand-alone application or can be combined with other modules into an enterprise application.

The Application Assembly Tool (AAT) provides flexibility in assembling Web modules. Options described below include:

- Importing an existing Web module (WAR file)
- Creating a new Web module
- Copying code artifacts (such as servlets) from one Web module into a new Web module

Although you can input various properties for Web archives, available properties are specific to the Servlet, JSP, and J2EE specification level.

### Steps for this task

1. Start the AAT.
2. From the New tab, select **Web Module**. Click **OK**.

The navigation tree now displays various sets of properties for configuring the new Web module.
3. **(Optional)** Use the property dialog shown in the AAT workspace to change the default file name and location.
  - a. It is recommended that you change the display name so that it differs from the file name.
  - b. If you like, change the temporary location of the Web module from the default location, `install_root/bin`.
4. Add at least one Web component (servlet or JSP file) to the module.

You must add at least one Web component, using one of the following methods.

  - Import an existing WAR file containing Web components.
    - a. In the navigation tree, right-click the **Web Components** folder.
    - b. Select **Import** from its right-click menu.
    - c. Use the file browser to locate and select the archive file for the module.
    - d. Click **Open**. The Web applications in the selected archive are displayed.
    - e. Select a Web application. Its Web components are displayed in the workspace.
    - f. Select the servlets or JSP files to be added and click **Add**. The components are displayed in the **Selected Components** window.
    - g. Click **OK**. The properties associated with the archive are also imported. The property dialog boxes in the workspace are populated automatically with values.
    - h. Double-click the Web Components icon to verify that the servlets or JSP files are included in the module.
    - i. Double-click the **Web Components** icon to verify that the servlets or JSP files are included in the module.
    - j. Save the Web module.
  - Copy and paste archive files from an existing module into the new Web component.
  - Create a new Web component.
    - a. In the navigation tree, right-click the **Web Components** folder.
    - b. Select **New** from its right-click menu.
    - c. When the new module is displayed, enter a component name and choose a component type.
    - d. Use the file browser to locate and select the archive file for the module.
    - e. Click the plus sign (+) to verify its contents and enter assembly properties.
    - f. In the **New Web Component** property dialog box, click **OK**.
    - g. Verify that the Web component has been added to the module by double-clicking the **Web components** icon in the navigation tree.
    - h. Click the component to view its corresponding property dialog box in the bottom portion of the pane.
5. Enter assembly properties for each Web component.

- a. Click the plus sign (+) next to the each component to reveal its property groups.
  - b. Right-click each property group icon and click **New** to display properties in the workspace.
6. Specify additional properties for the Web module.  
Right-click each property group's icon. Choose **New** to add new values, or edit existing values in the property pane. (Click **Help** for descriptions of the settings).  
Note that if you add a security constraint, you must add at least one Web resource collection.
7. Add any other files needed by the application.
- In the navigation tree, click the plus sign (+) next to the Files icon. Right-click **Add Class Files**, **Add JAR Files**, or **Add Resource Files**. Select **Add Files**.
  - Add files, using the Add Files dialog.
8. Save the application

### What to do next

Assemble other new modules of your choice, if needed:

- Assembling EJB modules.
- Assembling application client modules.
- "Assembling resource adapter modules" (see WebSphere Application Server documentation on resources).

You can also migrate existing modules.

Another option is to proceed directly to assembling a new application module. While assembling an application module, you can create any new modules that you need.

## Context parameters

A servlet context defines a server's view of the Web application within which the servlet is running. The context also allows a servlet to access resources available to it.

Using the context, a servlet can log events, obtain URL references to resources, and set and store attributes that other servlets in the context can use. These properties declare a Web application's parameters for its context. They convey setup information, such as a webmaster's e-mail address or the name of a system that holds critical data.

## Security constraints

Security constraints declare how Web content is to be protected.

Security constraints declare how to protect Web content. These properties associate security constraints with one or more Web resource collections. A constraint consists of a Web resource collection, an authorization constraint and a user data constraint.

- A Web resource collection is a set of resources (URL patterns) and HTTP methods on those resources. All requests that contain a request path that



matches the URL pattern described in the Web resource collection is subject to the constraint. If no HTTP methods are specified, then the security constraint applies to all HTTP methods.

- An authorization constraint is a set of roles that users must be granted in order to access the resources described by the Web resource collection. If a user who requests access to a specified URI is not granted at least one of the roles specified in the authorization constraint, the user is denied access to that resource.
- A user data constraint indicates that the transport layer of the client or server communications process must satisfy the requirement of either guaranteeing content integrity (preventing tampering in transit) or guaranteeing confidentiality (preventing reading while in transit).

## Servlet mappings

A servlet mapping is a correspondence between a client request and a servlet.

Servlet containers use URL paths to map client requests to servlets, and follow the URL path-mapping rules as specified in the Java Servlet specification. The container uses the URI from the request, minus the context path, as the path to map to a servlet. The container chooses the longest matching available context path from the list of Web applications that it hosts.

## Invoker attributes

Invoker attributes are used by the servlet that implements the invocation behavior.

## Error pages

Error page locations allow a servlet to find and serve a URI to a client based on a specified error status code or exception type.

These properties are used if the error handler is another servlet or JSP file. The properties specify a mapping between an error code or exception type and the path of a resource in the Web application. The container examines the list in the order that it is defined, and attempts to match the error condition by status code or by exception class. On the first successful match of the error condition, the container serves back the resource defined in the Location property.

## File serving

File serving allows a Web application to serve static file types, such as HTML. File-serving attributes are used by the servlet that implements file-serving behavior.

## Initialization parameters

Initialization parameters are sent to a servlet in its `HttpConfig` object when the servlet is first started.

## Servlet caching

Dynamic caching can be used to improve the performance of servlet and JavaServer Pages (JSP) files by serving requests from an in-memory cache. Cache entries contain the servlet's output, results of the servlet's execution, and metadata.

## Web components

A web component is a servlet, Java Server Page (JSP), or HTML file. One or more web components make up a web module.

## Web property extensions

Web property extensions are IBM extensions to the standard deployment descriptors for Web applications. These extensions include mime filtering and servlet caching.

## Web resource collections

A Web resource collection defines a set of URL patterns (resources) and HTTP methods belonging to the resource.

HTTP methods handle HTTP-based requests, such as GET, POST, PUT, and DELETE. A URL pattern is a partial Uniform Resource Locator that acts as a template for matching the pattern with existing full URLs in an attempt to find a valid file.

## Welcome files

A Welcome file is an entry point file (for example, `index.html`) for a group of related HTML files.

Welcome files are located by using a group of partial URIs. The Web container uses the partial URIs to find a valid file when the initial URI is not found.

## Context parameter assembly settings

A servlet context defines the server view of the Web application within which the servlet is running. The context also allows a servlet to access resources available to it. Using the context, a servlet can log events, obtain URL references to resources, and set and store attributes that other servlets in the context can use.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance*

### Parameter name (Required, String)

Specifies the name of a parameter, for example, `dataSourceName`.

Data type String

### Parameter value (Required, String)

Specifies the value of a parameter, for example, `jdbc/sample`.

Data type String

### Description

Contains a description of the context parameter.

Data type String

## Initialization parameter assembly settings

Use this page to specify the initialization parameters that are sent to a servlet in its `HttpConfig` object when the servlet is first started.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance* > **Web Components** > *component\_instance* > **Initialization Parameters**

### **Parameter name (Required, String)**

Specifies the name of an initialization parameter.

Data type String

### **Parameter value (Required, String)**

Specifies the value of the initialization parameter.

Data type String

### **Description**

Contains text describing the use of the parameter.

Data type String

## **Filter assembly settings**

Use the Filter panel to configure your filter settings.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance*

### **Filter name**

Specifies the logical name of the filter. This name maps the filter.

Data type String

### **Class**

Specifies the fully qualified classname of the filter.

Data type String

### **Description**

Provides a description of the filter.

Data type String

## **JavaServer Pages attribute assembly settings**

Use the JavaServer Pages (JSP) attributes page to set JSP attributes that are used by servlets that implement JSP processing behavior.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance* > **Assembly Property Extensions**

### **JSP Attribute (Name)**

Specifies the name of an attribute.

Data type String

## JSP Attribute (Value)

Specifies the value of an attribute.

Data type String

The WebSphere JSP container supports the following JSP attributes:

### **classdebuginfo**

Indicates the compiler should include debugging information in the generated classfile.

`classdebuginfo` `true` or `false`

Default is `false`.

### **classpath**

Specifies an additional classpath for compiling the generated servlets.

`classpath` `classpath` or `null`

Default is `null`.

### **deprecation**

Indicates the compiler should generate deprecation warnings when compiling the generated Java source.

`deprecation` `true` or `false`

Default is `false`.

### **disableJspRuntimeCompilation**

Indicates the runtime behavior of the JSP compiler. If this option is set to `true`, the JSP compiler does not compile or translate the JSP files, and the JSP engine only loads precompiled classfiles.

`disableJspRuntimeCompilation` `true` or `false`

Default is `false`.

### **ieClassID**

Indicates the Java plugin COM class ID for Internet Explorer. The `<jsp:plugin>` tags use this value.

`ieClassID` `classid`

Default is `clsid:8AD9C840-044E-11D1-B3E9-00805F499D93`.

### **javaEncoding**

Indicates the Java platform encoding to use to generate the JSP page servlet.

`javaEncoding` `encoding value`

Default is UTF-8.

### **jspCompilerPath**

Indicates the path of the compiler to use for compiling JSP pages.

jspCompilerPath *path name or null*

Default is *null*.

### **keepgenerated**

Indicates the Java files generated by the JSP compiler during the translation phase of the processing should be kept.

keepgenerated true or false

Default is false.

### **largefile**

Specifies support for large files. When the Java code is generated, the HTML data in a JSP file is stored separately instead of being saved as constant string data in the generated servlet.

largefile true or false

Default is false.

### **mappedfile**

Indicates the compiler should generate Java source that includes a print statement for every line in the JSP file. Use this option for debugging purposes only. It is not recommended for production environments because the mappedfile option generates too many `out.print()` statements.

mappedfile true or false

Default is false.

### **scratchdir**

Specifies the directory where the generated classfiles are created.

scratchdir *directory name.*

Default is `[WAS_INSTALL_ROOT]/temp`.

**Note:** The system property `com.ibm.websphere.servlet.temp.dir` can be used to set the `scratchdir` option on a server-wide basis. This setting, if it is present, overrides the system property.

### **usePageTagPool**

Enables or disables the reuse of custom tag handlers on an individual JavaServer Page basis.

usePageTagPool true or false

Default is false.

The note in the `useThreadTagPool` attribute description also applies to `theusePageTagPool` attribute.

### **useThreadTagPool**

Enables or disables the reuse of custom tag handlers on a per request thread basis.

`useThreadTagPool` true or false

Default is false.

The note in the `useThreadTagPool` attribute description also applies to `theusePageTagPool` attribute.

**Note:** Enabling custom tag handler reuse might reveal problems in your tag handler code regarding the tags ability to be reused. A custom tag handler should always do two things:

1. The `release()` method of the tag handler should reset its state and release any private resources that it might have used. The JSP engine guarantees the `release()` method will be called before the tag handler is garbage collected.
2. In the `doEndTag()` method, all instance states associated with this instance must be reset.

### **verbose**

Indicates the compiler should generate verbose output when compiling the generated Java source code.

`verbose` true or false

Default is false.

## **Multipurpose Internet Mail Extensions (MIME) filter assembly settings**

Use this page to configure Multipurpose Internet Mail Extensions (MIME) filters.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance* > **Assembly Property Extensions**

### **Component name (Required, String)**

Specifies the name of the servlet or JavaServer Pages(TM) (JSP) file. This name must be unique within the Web module.

Data type String

### **Display name**

Specifies a short name that is intended for display by GUIs.

Data type String

### **Description**

Contains a description of the servlet or JSP file.

Data type String

### Component type

Specifies the type of Web component. Valid values are `servlet` or `JSP file`.

Data type String

### Class name (Required, String)

Specifies the full path name for the servlet class.

Data type String

### JSP file (Required, String)

Specifies the full path name for the JSP file.

Data type String

### Load on startup

Indicates whether this servlet loads at the startup of the Web application.

The default is `false` (the check box is not selected). Also specifies a positive integer indicating the order in which to load the servlet. Lower integers are loaded before higher integers. If no value is specified, or if the value specified is not a positive integer, the container is free to load the servlet at any time in the startup sequence.

Data type String

### Small icon

Specifies a JPEG or GIF file containing a small image (16x16 pixels). Use the image as an icon to represent the Web component in a GUI.

Data type JPEG, GIF

### Large icon

Specifies a JPEG or GIF file containing a large image (32x32 pixels). Use the image as an icon to represent the Web component in a GUI.

Data type JPEG, GIF

## Page list assembly settings

Page lists allow you to avoid hardcoding URLs in servlets and JSP files.

Access this page by traversing the following path in the Application AssemblyTool:

**Web Modules** > *component\_instance* > **Web Components**

### Name

Specifies the name of the markup language—for example, Hypertext Markup Language (HTML), Wireless Markup Language (WML), and Voice Extensible Markup Language (VXML).

Data type String

### MIME Type

Specifies the Multi-Purpose Internet Mail Extensions (MIME) type of the markup language, for example, `text/html` and `text/x-vxml`.

Data type String

### **Error Page**

Specifies the name of an error page.

Data type String

### **Default Page**

Specifies the name of a default page.

Data type String

### **Pages - Name**

Specifies the name of the page to serve, for example, *StockQuoteRequest.page*.

Data type String

### **Pages - URI**

Specifies the URI of the page to serve, for example, *examples/StockQuoteHTMLRequest.jsp*.

Data type String

## **Security constraint assembly settings**

Use the Security constraints panel to configure security constraints.

To view this Application Assembly Tool (AAT) panel, open an existing or create a new Web module. Right-click **Security Constraints** from the left navigation menu. Click **New**.

If multiple security constraints are specified, the container uses the "first match wins" rule when processing a request to determine what authentication method to use, or what authorization to allow

### **Security constraint name**

Specifies the name of the security constraint.

Data type String

### **Authorization Constraints - Roles**

Specifies the user roles that are permitted access to this resource collection.

Data type String

### **Authorization Constraints - Description**

Contains a description of the authorization constraints

Data type String

### **User Data Constraints - Transport guarantee**

Indicates how data communicated between the client and the server is to be protected.



Specifies that the protection for communications between the client and server is None, Integral, or Confidential.

- None means that the application does not require any transport guarantees.
- Integral means that the application requires that the data sent between the client and the server must be sent in such a way that it cannot be changed in transit.
- Confidential means that the application requires that the data must be transmitted in a way that prevents other entities from observing the contents of the transmission.

In most cases, Integral or Confidential indicates that the use of SSL is required.

Data type String

### **User Data Constraints - Description**

Contains a description of the user data constraints.

Data type String

## **Servlet mapping assembly settings**

A servlet mapping is a correspondence between a client request and a servlet.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance*

### **URL pattern (Required, String)**

Specifies the URL pattern of the mapping.

The URL pattern must conform to the Servlet specification. Use the following syntax:

- A string beginning with a slash character (/) and ending with the slash and asterisk characters (/\*) represents a path mapping.
- A string beginning with the characters \*. represents an extension mapping.
- All other strings are used as exact matches only.
- A string containing only the slash character (/) indicates that the servlet specified by the mapping becomes the default servlet of the application. In this case, the servlet path is the request Uniform Resource Identifier (URI) minus the context path, and the path information is null.

Data type String

### **Servlet (Required, String)**

Specifies the name of the servlet associated with the URL pattern.

Data type String

## **Tag library assembly settings**

Use this page to define the tag library parameters.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance*

Java ServerPages (JSP) tag libraries contain classes for common tasks such as processing forms and accessing databases from JSP files.

### **Tag library file name (Required, String)**

Specifies a file name relative to the location of the web.xml document, identifying a tag library used in the Web application.

Data type String

### **Tag library location (Required, String)**

Contains the location, as a resource relative to the root of the Web application, where you can find the Tag Library Definition file for the tag library.

Data type String

## **Welcome file assembly settings**

Use this page to configure your welcome page.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance*

### **Welcome file (Required, String)**

The Welcome file list is an ordered list of partial URLs with no trailing or leading slash characters (/).

The Web server appends each file in the order specified and checks whether a resource in the Web archive (WAR) file is mapped to that request Uniform Resource Identifier (URI). The container forwards the request to the first resource in the WAR file that matches.

Data type String

## **Servlet caching configuration assembly settings**

Use this page to configure your cache groups. Access this page by traversing the following path in the Application AssemblyTool: **Web Modules** >

*component\_instance* > **Assembly Property Extensions**

The properties on the General tab define a cache group and govern how long an entry remains in the cache. The properties on the ID Generation tab define how cache IDs are built and the criteria used to cache or invalidate entries. The properties on the Advanced tab define external cache groups and specify custom interfaces for handling servlet caching.

### **Caching group name**

Specifies a name for the group of servlets or JavaServer Pages (JSP) files to cache.

### **Priority**

Defines the default priority for cached servlets. Specify as an integer. The default value is 1.

Priority is an extension of the Least Recently Used (LRU) caching algorithm. It represents the number of cycles through the LRU algorithm that an entry is guaranteed to stay in the cache. The priority represents the length of time that an entry remains in the cache before becoming eligible for removal. On each cycle of the algorithm, the priority of an entry is decremented. When the priority reaches zero, the entry is eligible for invalidation. If an entry is requested while in the cache, its priority is reset to the priority value. Regardless of the priority value and the number of requests, an entry is invalidated when its timeout occurs. Consider increasing the priority of a servlet or JSP file when it is difficult to calculate the output of the servlet or JSP file or when the servlet or JSP file is executed more often than average. Priority values should be low. Higher values do not yield much improvement but use extra LRU cycles. Use timeout to guarantee the validity of an entry. Use priority to rank the relative importance of one entry to other entries. Giving all entries equal priority results in a standard LRU cache that increases performance significantly.

### **Timeout**

Specifies the length of time, in seconds, that a created cache entry remains in the cache.

When this time elapses, the entry is removed from the cache. If the timeout is zero or a negative number, the entry does not time out. It is removed when the cache is full or programmatically, from within an application.

### **Invalidate only**

Specifies that invalidations for a servlet take place, but that no servlet caching is performed.

For example, you can use this property to prevent caching of control servlets. Control servlets treat HTTP requests as commands and execute those commands. By default, this check box is not selected.

### **Caching group members**

Specifies the names of the servlets or JSP files to cache. The URIs are determined from the servlet mappings.

### **Use URIs for cache ID building**

Specifies whether or not to use the URI of the requested servlet to create a cache ID. By default, URIs are used.

### **Use specified string**

Specifies a string representing a combination of request and session variables to use for creating cache IDs. This property defines request and session variables, and the cache uses the values of these variables to create IDs for the entries.

### **Variables - ID**

Specifies the name of a request parameter, request attribute, session parameter, or cookie.

### **Variables - Type**

Specifies the type of variable indicated in the ID field. The valid values are Request parameter, Request attribute, Session parameter, or Cookie.

### **Variables - Method**

Specifies the name of a method in the request attribute or session parameter. The output of this method is used to generate cache entry IDs. If this value is not specified, the toString method is used by default.

### **Variables - Data ID**

Specifies a string that, combined with the value of the variable, generates a group name for the cache entry. The cache entry is placed in this group. You can invalidate this group.

### **Variables - Invalidate ID**

Specifies a string that is combined with the value of the variable on the request or session to form a group name. The cache invalidates the group name.

### **Required**

Specifies whether a value must exist in the request. If this check box is selected, and either the request parameter, request attribute, session parameter, or the method is not specified, the request is not cached.

### **External cache groups - Group name**

Specifies the name of the external cache group to which this servlet is published.

### **ID generator**

Specifies a user-written interface for handling parameters, attributes, and sessions.

The value must represent a full package and class name of a class extending `com.ibm.websphere.servlet.cache.IdGenerator`. The properties specified in the Application Assembly Tool are used and passed to the `IdGenerator` in the `initialize` method inside a `com.ibm.websphere.servlet.cache.CacheConfig` object.

Data type String

### **Meta data generator**

Specifies a user-written interface for handling invalidation, priority levels, and external cache groups.

The value must represent the full package and class name of a class extending `com.ibm.websphere.servlet.cache.MetaDataGenerator`. The properties specified in the Application Assembly Tool are used and passed to the `MetaDataGenerator` in the `initialize` method inside a `com.ibm.websphere.servlet.cache.CacheConfig` object.

Data type String

## **Web components assembly settings**

Use this page to set the assembly properties for the components that make up a Web module.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance* > **Web Components**

### **Component name**

Specifies the name of the servlet or JavaServer Pages(TM) (JSP) file. This name must be unique within the Web module.

Data type String

### **Display name**

Specifies a short name that is intended for display by GUIs.

Data type String

### Description

Contains a description of the servlet or JSP file.

Data type String

### Component type

Specifies the type of Web component. Valid values are servlet or JSP file.

Data type String

### Class name

Specifies the full path name for the servlet class.

Data type String

### JSP file

Specifies the full path name for the JSP file.

Data type String

### Load on startup

Indicates whether this servlet loads at the startup of the Web application. The default is false (the check box is not selected).

This field also specifies a positive integer indicating the order in which the servlet is to load. Lower integers are loaded before higher integers. If no value is specified, or if the value specified is not a positive integer, the container is free to load the servlet at any time in the startup sequence.

Data type Boolean  
Default False

### Small icon

Specifies a JPEG or GIF file containing a small image (16x16 pixels). Use the image as an icon to represent the Web component in a GUI.

Data type String

### Large icon

Specifies a JPEG or GIF file containing a large image (32x32 pixels). Use the image as an icon to represent the Web component in a GUI.

Data type String

### Run as role name

Enter a role name that represents the user account under which the servlet executes. The default role name is blank, which indicates the servlet runs under the user that logged into the application server. The role name of "all role" indicates the servlet can execute under different users.

Data type String

## Description

In this optional field, enter a description that explains the importance of the role, and where and how the role can be used.

Data type String

## Run as role mode

Indicates a security role that is defined in the enterprise application.

Data type String

## Local Transactions - Unresolved action

Specifies the action the Web container must take if resources in a local transaction are uncommitted by an application. This property is an IBM extension to the standard J2EE deployment descriptor. A local transaction context is provided by the container in the absence of a global transaction context.

|                        |                                      |
|------------------------|--------------------------------------|
| Data type              | String                               |
| Default                | Rollback                             |
| Range valid values are | Commit                      Rollback |

Additional information about these settings follows:

### Commit

At end of the local transaction context, the container instructs all the unresolved local transactions to commit.

### Rollback (default)

At end of the local transaction context, the container instructs all the unresolved local transactions to rollback.

## WebSphere Application Server Enterprise assembly settings for Web components

Use this page to configure Enterprise functions for Web components (servlets and JSPs).

**Own task:** A J2EE component's own task defines the task that can be used to identify a unit of work begun by this component to application profile-configured policies on downstream requests.

**Name:** The name of the task to be propagated on requests from this J2EE component.

Datatype String

**Description:** The description of the task.

An optional field provided as a convenience to the developer.

Datatype String

**Task references:** Task references enable developers to programmatically set the current task under which a component is executing.

**Name:** The name of the task reference. This name corresponds to the string that the developer uses to set the current task.

**Task:** The task that the run-time environment will associate with execution of the current component.

**Name:** The name of the task.

Programmatically setting the logical task name causes this task name to be associated with requests from this J2EE component.

**Description:** A description of the task.

An optional field provided as a convenience for the developer.

**ActivitySession control kind:** The action to be taken by the TransactionControl collaborator that runs in the web container.

Default

None

Range

### **Application**

Access to UserActivitySessions is provided. Any HttpSession started by a servlet does not have an ActivitySession automatically associated with it by the container. The container associates any ActivitySession started by the servlet with an HttpSession already available to the servlet. An ActivitySession started before any HttpSession is created is not associated with any HttpSession and must be completed by the application before dispatch-end; otherwise it is cleaned up by the container. Servlets with this attribute cannot be included or forwarded to.

### **Container**

A servlet has no access to UserActivitySessions. Any HttpSession started by the servlet has an ActivitySession automatically associated with it by the container, and this ActivitySession is put onto the thread of execution. If such a servlet is dispatched by a request that has an HttpSession containing no ActivitySession, then the container starts an ActivitySession and associates it with the HttpSession and the thread.

### **None**

A servlet has no access to UserActivitySessions, and no participation in an ActivitySession is tolerated. Any HttpSession containing an ActivitySession that is associated with a request dispatched on such a servlet is rejected with a ServletException.



**Internationalization type:** The Internationalization type of a servlet indicates whether the servlet or the hosting J2EE Web container will manage internationalization context on lifecycle method invocations. For servlets, the Internationalization type setting can be configured to Application or Container, and defaults to Container. The setting applies to all lifecycle methods of a servlet.

Select whether the servlet will employ **Application** or **Container** internationalization context management using the **Internationalization type** drop-down menu.

Default

Container

Range

### Application

Under Application-managed Internationalization (AMI) servlets are responsible to manage (set) invocation context elements using the internationalization context API. The container suspends internationalization context imported from the calling process on lifecycle method invocations, including HTTP servlet requests; to continue propagating this context on subsequent business method requests, the method must use the API to transfer caller context elements to the invocation context scoped to the method. Invocation context elements not set using the API default to the respective elements of the current process when accessed using the API or when propagated on outgoing business method requests.

Specify Application-managed Internationalization for servlets having internationalization context management requirements not supported by Container-managed Internationalization.

### Container

Under Container-managed Internationalization (CMI) servlets have read-only access to the internationalization context API and rely solely on the Web container to manage invocation context in accordance with the **Run as** field of the applicable Container Internationalization attribute. On a service method invocation, the container scopes the context indicated by the **Run as** field to the method; when the method has completed, the container removes this context from scope. If the servlet is not indicated within a Container Internationalization attribute, the container scopes the caller context by default. Invocation context elements unavailable to the container default to the respective elements of the server JVM when accessed using the API or when propagated on outgoing business method requests.

## Web modules assembly settings

Use this page to set the assembly properties for web modules. Web modules are composed of one or more web components.

Access this page by traversing the following path in the Application AssemblyTool:  
*application\_instance* > **Web Modules**

**File name**

Specifies the file name of the Web module, relative to the top level of the application package.

**Alternative DD**

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module's JAR file.

This file is the postassembly version of the deployment descriptor file. (The original deployment descriptor file can be edited to resolve dependencies and security information. Directing the use of the alternative deployment descriptor allows you to keep the original deployment descriptor file intact). The value of the Alternative DD property must be the full path name of the deployment descriptor file relative to the module's root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read directly from the module's JAR file.

**Context root**

Specifies the context root of the Web application. The context root is combined with the defined servlet mapping (from the WAR file) to compose the full URL that users type to access the servlet.

For example, if the context root is /gettingstarted and the servlet mapping is MySession, then the URL is `http://host:port/gettingstarted/MySession`.

**Classpath**

Specifies the class path for resources used by the Web application, relative to the ear file..

If your Web application requires access to classes within an ear file, specify the relative path of the classes in this field.

**Display name**

Specifies a short name that is intended to be displayed by GUIs.

**Description**

Contains a description of the Web module.

**Distributable**

Specifies that this Web application is programmed appropriately to deploy into a distributed servlet container.

**Small icon**

Specifies a JPEG or GIF file containing a small image (16x16 pixels). The image is used as an icon to represent the module in a GUI.

**Large icon**

Specifies a JPEG or GIF file containing a large image (32x32 pixels). The image is used as an icon to represent the module in a GUI.

**Session configuration**

Indicates that session configuration information is present. Checking this box makes the Session timeout property editable.

**Session timeout**

Specifies a time period, in seconds, after which a client is considered inactive. The default value is zero, indicating that the session timeout never expires.

## **Login configuration — Authentication method**

Specifies an authentication method to use. As a prerequisite to gaining access to any Web resources protected by an authorization constraint, a user must authenticate by using the configured mechanism.

A Web application can authenticate a user to a Web server by using one of the following mechanisms: HTTP basic authentication, HTTP digest authentication, HTTPS client authentication, and form-based authentication.

- HTTP basic authentication is not a secure protocol because the user password is transmitted with a simple Base64 encoding and the target server is not authenticated. In basic authentication, the Web server requests a Web client to authenticate the user and passes a string called the realm of the request in which the user is to be authenticated.
- HTTP digest authentication transmits the password in encrypted form.
- HTTPS client authentication uses HTTPS (HTTP over SSL) and requires the user to possess a public key certificate.
- Form-based authentication allows the developer to control the appearance of login screens.

The Login configuration properties are used to configure the authentication method that should be used, the realm name that should be used for HTTP basic authentication, and the attributes that are needed by the form-based login mechanism. Valid values for this property are Unspecified, Basic, Digest, Form, and Client certification.

### **Note:**

HTTP digest authentication is not supported as a login configuration in this product. Also, not all login configurations are supported in all of the product's global security authentication mechanisms (Local Operating system, LTPA, and custom pluggable user registry). HTTP basic authentication and form-based login authentication are the only authentication methods supported by the Local Operating system user registry. LTPA and the custom pluggable user registry are capable of supporting HTTP basic authentication, form-based login, and HTTPS client authentication.

## **Login configuration — Realm name**

Specifies the realm name to use in HTTP basic authorization. It is based on a user name and password, sent as a string (with a simple Base64 encoding).

An HTTP realm is a string that allows URIs to be grouped together. For example, if a user accesses a secured resource on a Web server within the "finance realm," subsequent access to the same or different resource within the same realm does not result in a repeat prompt for a user ID and password.

## **Login configuration — Login page**

Specifies the location of the login form. If form-based authentication is not used, this property is disabled.

## **Form Login Config — Error page**

Specifies the location of the error page. If form-based authentication is not used, this property is disabled.

## **Reload interval**

Specifies a time interval, in seconds, in which the file system of the Web application is scanned for updated files. The default is 3 seconds.

### **Reloading enabled**

Specifies whether file reloading is enabled. The default is true.

### **Default error page**

Specifies a file name for the default error page. If no other error page is specified in the application, this error page is used.

### **Additional classpath**

Specifies the full class path that will be used to reference classes outside of those specified in the archive.

If your Web application requires access to classes not contained in the archive file, specify the full path for those classes in this field.

### **File serving enabled**

Specifies whether file serving is enabled. File serving allows the application to serve static file types, such as HTML and GIF. File serving can be disabled if the application contains only dynamic components. The default value is true.

### **Directory browsing enabled**

Specifies whether directory browsing is enabled. Directory browsing allows the application to browse disk directories. Directory browsing can be disabled if, you want to protect data. The default value is true.

### **Serve servlets by classname**

Specifies whether a servlet can be served by requesting its class name.

Usually, servlets are served only through a URI reference. The class name is the actual name of the servlet on disk. For example, a file named SnoopServlet.java compiles into SnoopServlet.class. (This is the class name.) SnoopServlet.class is normally invoked by specifying snoop in the URI. However, if Serve Servlets by Classname is enabled, the servlet is invoked by specifying SnoopServlet. The default value is true.

### **Virtual hostname**

Specifies a virtual host name. A virtual host is a configuration enabling a single host machine to resemble multiple host machines. This property allows you to bind the application to a virtual host in order to enable execution on that virtual host.

### **Filter mappings**

Specifies the filter mapping declarations in this application. The container uses the filter mapping declarations to decide on the type and order of filters to apply to a request.

After the container matches the request URI to a servlet, for each filter mapping element, it determines what filters to apply based on the servlet name or the URL pattern, depending on the style specified. Filters are invoked in the same order as the one specified in the list of filter mapping elements. The value that you specify for the filter name must be the same value as that specified in the `<filter><filtername>` sub-element declarations in the deployment descriptor.

## **Assembly property extensions**

Use this panel to configure WebSphere Application Server specific Web module extensions, or also referred to as assembly property extensions.

This panel lists the extensions that can be configured through the tool.

Reach the applicable extension panel by clicking on the panel name in the navigation at the left, or by double-clicking the attribute name in the list provided.

Access this page by traversing the following path in the Application AssemblyTool:  
**Web Modules** > *component\_instance* > **Assembly Property Extensions**

Assembly properties for a Web module include:

- File serving attributes
- Invoker attributes
- JavaServer Pages (JSP) attributes
- Multipurpose Internet Mail Extensions (MIME) filters
- Servlet caching configurations

## File serving attribute assembly settings

File serving allows a Web application to serve static file types, such as HTML. File-serving attributes are used by the servlet that implements file-serving behavior.

Access this page by traversing the following path in the Application AssemblyTool:  
**Web Modules** > *component\_instance* > **Assembly Property Extensions**

### File Serving Attribute (Name)

Specifies the name of an attribute.

Data type String

### File Serving Attribute (Value)

Specifies the value of an attribute.

Data type String

## Invoker attribute assembly settings

Invoker attributes are used by the servlet that implements the invocation behavior.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance* > **Assembly Property Extensions**

### Invoker Attribute (Name)

Specifies the name of an attribute.

Data type String

### Invoker Attribute (Value)

Specifies the value of an attribute.

Data type String

## Error page assembly settings

Error page locations allow a servlet to find and serve a URI to a client based on a specified error status code or exception type. These properties are used if the error handler is another servlet or JSP file.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance*

The error page properties specify a mapping between an error code or exception type and the path of a resource in the Web application. The container examines the list in the order that it is defined, and attempts to match the error condition by status code or by exception class. On the first successful match of the error condition, the container serves back the resource defined in the Location property.

### **Error code**

Indicates that the error condition is a status code.

Data type Integer

### **Error Code (Required, String)**

Specifies an HTTP error code, for example, 404.

Data type String

### **Exception**

Indicates that the error condition is an exception type.

Data type String

### **Exception type name (Required, String)**

Specifies an exception type.

Data type String

### **Location (Required, String)**

Contains the location of the error-handling resource in the Web application.

Data type String

## **Web resource collections security constraint properties**

A Web resource collection defines a set of URL patterns or resources and HTTP methods belonging to the resource, which define the security constraints for a Web component.

Access this page by traversing the following path in the Application Assembly Tool:

**Web Modules** > *component\_instance*

HTTP methods handle HTTP-based requests, such as GET, POST, PUT, and DELETE. A URL pattern is a partial Uniform Resource Locator that acts as a template for matching the pattern with existing full URLs in an attempt to find a valid file.

### **Web resource name**

Specifies the name of a Web resource collection.

Data type String

## Web resource description

Contains a description of the Web resource collection.

### HTTP methods

Specifies the HTTP methods to which the security constraints apply. If no HTTP methods are specified, then the security constraint applies to all HTTP methods. The valid values are GET, POST, PUT, DELETE, HEAD, OPTIONS, and TRACE.

Data type String

### URL pattern

Specifies URL patterns for resources in a Web application. All requests that contain a request path that matches the URL pattern are subject to the security constraint.

Data type String

---

## Troubleshooting tips for Web application deployment

Deployment of a Web application is successful if you can access the application by typing a Uniform Resource Locator (URL) in a browser, or if you can access the application by following a link.

If you cannot access your application, follow these steps to eliminate some common errors that can occur during migration or deployment.

### Web module does not run in WebSphere Application Server Version 5.

|                      |  |
|----------------------|--|
| Symptom              | Your Web module does not run when you migrate it to Version 5  |
| Problem              | In Version 4.x, the classpath setting that affected visibility was <i>Module Visibility Mode</i> . In Version 5, you must use class loader policies to set visibility.   |
| Recommended response | Reassemble an existing module, or change the visibility settings in the class loader policies. in the class loader policies.<br><br>See article (Migration of module visibility modes from Version 4.x) for more information and examples. |

### Welcome page is not visible.

|                      |   |
|----------------------|---|
| Symptom              | You cannot access an application with a Web path of:<br><code>/webapp/myapp</code>  |
| Problem              | The default welcome page for a Web application is assumed to be <i>index.html</i> . You cannot access the default page of the <i>myapp</i> application unless it is named <i>index.html</i> . |
| Recommended response | To identify a different welcome page, modify the properties of the Web module during assembly. See article Assembling Web modules for more information.                                       |

### HTML files are not found.



|                      |   |
|----------------------|---|
| Symptom              | Your Web application ran successfully on prior versions, but now you encounter errors that the welcome page (typically <i>index.html</i> ), or referenced HTML files are not found:<br>Error 404: File not found: Banner.html<br>Error 404: File not found: HomeContent.html  |
| Problem              | For security and consistency reasons, Web application URLs are now case-sensitive on all operating systems.<br><br>Suppose the content of the index page is as follows:<br><!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 5.0 Frameset//EN"><br><HTML><br><TITLE><br>Insurance Home Page<br></TITLE><br><frameset rows="18,80"><br><frame src="Banner.html" name="BannerFrame" SCROLLING=N<br><frame src="HomeContent.html" name="HomeContentFrame"><br></frameset><br></HTML> |
| Recommended response | However the actual file names in the<br>\WebSphere\AppServer\installedApps\... directory where<br>the application is deployed are:<br>banner.html<br>homecontent.html<br>To correct this problem, modify the <i>index.html</i> file to change<br>the names <i>Banner.html</i> and <i>HomeContent.html</i> to <i>banner.html</i><br>and <i>homecontent.html</i> to match the names of the files in the<br>deployed application.  |

---

## Modifying the default Web container configuration

### Before you begin

If you choose not to use the default application server, create an application server instance.

An application server instance contains a single instance of a Web container, which is created automatically when an application server is created.

The Web container is created initially with default properties values suitable for simple Web applications. However, these values might not be appropriate for more complex Web applications.

Your application is complex if it requires any of the following features:

- virtual host
- servlet caching
- special client request loads
- persistent HTTP session support
- special HTTP transport settings

Modify the following properties if you have a complex application:

### Steps for this task

1. **(Optional)** Modify the Web container General Properties if your Web application requires a virtual host, other than the default\_host, or requires servlet caching.
2. **(Optional)** Modify the Web Container Additional Properties Thread Pool setting if your application handles special client request loads.
3. **(Optional)** Modify the Web Container Additional Properties Session Management setting if your application requires persistent HTTP session support.
4. **(Optional)** Modify the Web Container Additional Properties HTTP transports setting if your application requires one of the following HTTP transport settings:
  - Unique hostname and port for client access
  - SSL enablement
5. **(Optional)** Modify the Web Container Additional Properties Custom Properties setting if your application requires global settings for internal servlets for WAR files packaged by third party tools.

## Web container

A Web container handles requests for servlets, JavaServer Pages (JSP) files, and other types of files that include server-side code. The Web container creates servlet instances, loads and unloads servlets, creates and manages request and response objects, and performs other servlet management tasks.

The Web server plug-ins, provided by the WebSphere Application Server, help supported Web servers pass servlet requests to Web containers.

## Web container settings

Use this page to configure the container settings.

Access this page by traversing the following path in the administrative console:

**Servers > Application Servers > *server\_instance* > Web container**

### Configuration - General Properties

#### Default virtual host

Specifies a virtual host is a configuration enabling a single host machine to resemble multiple host machines. Resources associated with one virtual host cannot share data with resources associated with another virtual host, even if the virtual hosts share the same physical machine.

Select a virtual host option:

#### Default Host

The product provides a default virtual host with some common aliases, such as the machine IP address, short host name, and fully qualified host name. The alias comprises the first part of the path for accessing a resource such as a servlet. For example, it is localhost:9080 in the request `http://localhost:9080/myServlet`.

#### Admin Host

This is another name for the application server; also known as *server1* in the base installation. This process supports the use of the administrative console.

## Servlet caching

Specifies that if a servlet is invoked once and it generates output to be cached, a cache entry is created containing not only the output, but also side effects of the invocation. These side effects can include calls to other servlets or Java Server Pages (JSP) files, as well as metadata about the entry, including timeout and entry priority information.

### Enable servlet caching

Check this box to enable servlet caching.

## Web module settings

Use this page to configure Web module settings.

Access this page by traversing the following path in the administrative console:

**Applications** > **Enterprise Application** > *application\_instance* > **Web Module**

### URI

Specifies a URI that, when resolved relative to the application URL, specifies the location of the module archive contents on a file system. The URI must match the ModuleRef URI in the deployment descriptor of an application if the module was packaged as part of a deployed application or enterprise archive (EAR) file.

### Name

Specifies the unique display name for the module.

### Alternate DD

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module JAR file.

This file is the *post-assembly* version of the deployment descriptor file. You can edit the original deployment descriptor file to resolve dependencies and security information. Specifying the use of the alternative deployment descriptor keeps the original deployment descriptor file intact.

The value of the *Alternate DD* property must be the full path name of the deployment descriptor file, relative to the module root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read from the module JAR file.

### Starting weight

Specifies the order in which modules are started. Lower weighted modules are started before higher weighted modules.

### Prefer WEB-INF Classes

Specifies classes to load in WEB-INF before any other classes. Implementing the application class loader is recommended so that classes and resources packaged within the WAR file load before classes and resources residing in container-wide library JAR files.

### Initial State

Specifies the default state of this application at server startup.

## Web Module Deployment settings

Use this page to configure an instance of Web module deployment.

Access this page by traversing the following path in the administrative console:

**Applications** > **Enterprise Application** > *application\_instance* > **Web Module** > *Web Module\_instance*

## URI

Specifies a URI that, when resolved relative to the application URL, specifies the location of the module archive contents on a file system. The URI must match the ModuleRef URI in the deployment descriptor of an application if the module was packaged as part of a deployed application or enterprise archive (EAR) file.

## Alternate DD

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module JAR file.

This file is the *post-assembly* version of the deployment descriptor file. You can edit the original deployment descriptor file to resolve dependencies and security information. Specifying the use of the alternative deployment descriptor keeps the original deployment descriptor file intact.

The value of the *Alternate DD* property must be the full path name of the deployment descriptor file, relative to the module root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read from the module JAR file.

## Starting weight

Specifies the order in which modules are started. Lower weighted modules are started before higher weighted modules.

## ClassLoader Mode

Specifies whether the class loader should search in the parent class loader or in the application class loader first to load a class. The standard for JDK class loaders and WebSphere class loaders is PARENT\_FIRST. By specifying PARENT\_LAST, your application can override classes contained in the parent class loader, but this action can potentially result in ClassCastException or LinkageErrors if you have mixed use of overridden classes and non-overridden classes.

The options are PARENT\_FIRST and PARENT\_LAST. The default is to search in the parent class loader before searching in the application class loader to load a class.

|           |              |
|-----------|--------------|
| Data type | String       |
| Default   | PARENT_FIRST |

## Web container custom property settings

Use this page to configure arbitrary name-value pairs of data, where the name is a property key and the value is a string value that can be used to set internal system configuration properties. Defining a new property enables you to configure a setting beyond that which is available in the administrative console.

To view this administrative console page, click **Servers** > **Application Servers** > *server\_name* > **Web Container** > > **Custom Properties** >

## Name

Specifies the name (or key) for the property.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

## Value

Specifies the value paired with the specified name.

Data type String

## Description

Provides information about the name-value pair.

Data type String

WAR files packaged using third party tools cannot specify behavior for the services exposed by the Web container internal servlets. You can globally enable/disable internal servlets for all Web applications at the Web container level by creating name/value pairs such as:

| Name                            | Value |
|---------------------------------|-------|
| fileServingEnabled              | true  |
| directoryBrowsingEnabled        | true  |
| serveServletsByClassnameEnabled | true  |

Settings defined at the Application AssemblyTool level will take precedence over the global settings set through the custom properties at the Web container level.

Web application deployment extensions will continue to hold configuration information for the services provided by the internal servlets, and will take precedence over the global settings set through the custom properties at the Web container level.

---

## Web applications: Resources for learning


Use the following links to find relevant supplemental information about Web applications. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming model and decisions
- Programming instructions and examples
- Programming specifications

### Programming model and decisions

-  [J2EE BluePrints for Web applications](#)
-  [Redbook on the design and implementation of Servlets, JSP files, and enterprise beans](#)

### Programming instructions and examples

- ✓ Redbook on Servlet and JSP file Programming
- 📖 Sun's Java™ Tutorial on Servlets
- 📖 Introduction to JavaServer Pages - Tutorial
- 📖 Bean Scripting Framework description
- 📖 Web delivered samples in the Samples Gallery

#### Programming specifications

- 📖 Java 2 Software Development Kit (SDK)
- 📖 Servlet 2.3 Specification
- 📖 JavaServer Pages 1.2 Specification
- 📖 Differences between JavaScript and ECMAScript
- 📖 ISO 8859 Specifications

---

## Chapter 3. Managing HTTP sessions

IBM WebSphere Application Server provides a service for managing HTTP sessions: Session Manager. The key activities for session management are summarized below.

Before you begin these steps, make sure you are familiar with the programming model for accessing HTTP session support in the applications following the Servlet 2.3 API.

### Steps for this task

1. Plan your approach to session management, which could include session tracking, session recovery, and session clustering.
2. Create or modify your own applications to use session support to maintain sessions on behalf of Web applications.
3. Assemble your application.
4. Deploy your application.
5. Ensure the administrator appropriately configures session management in the administrative domain.
6. Adjust configuration settings and perform other tuning activities for optimal use of sessions in your environment.

---

## Sessions

A session is a series of requests to a servlet, originating from the same user at the same browser.

Sessions allow applications running in a Web container to keep track of individual users.

For example, a servlet might use sessions to provide "shopping carts" to online shoppers. Suppose the servlet is designed to record the items each shopper indicates he or she wants to purchase from the Web site. It is important that the servlet be able to associate incoming requests with particular shoppers. Otherwise, the servlet might mistakenly add Shopper\_1's choices to the cart of Shopper\_2.

A servlet distinguishes users by their unique session IDs. The session ID arrives with each request. If the user's browser is cookie-enabled, the session ID is stored as a cookie. As an alternative, the session ID can be conveyed to the servlet by URL rewriting, in which the session ID is appended to the URL of the servlet or JavaServer Pages (JSP) file from which the user is making requests. For requests over HTTPS or Secure Sockets Layer (SSL), Another alternative is to use SSL information to identify the session.

---

## Migrating HTTP sessions

**Note:** In Version 5 default write frequency mode is TIME\_BASED\_WRITES, which is different from Version 4.0 and 3.5 default mode of END\_OF\_SERVICE.

### Migrating from Version 4.0

No programmatic changes are required to migrate from version 4.0 to version 5.

### Migrating from Version 3.5

If you have Version 3.5 applications running in Servlet 2.1 mode, some of the following Version 5 differences might influence how you choose to track and manage sessions.

#### Steps for this task

1. During application development, modify session-related APIs as needed.  
Some API changes are required in order to redeploy existing applications on Version 5. These include changes to the HttpSession API itself as well as issues associated with moving to support for the Servlet 2.3 specification. Certain Servlet 2.1 API methods have been deprecated in Servlet 2.3 API . These deprecated APIs still work in Version 5.0, but they may be removed in a future version of the API. Changes are summarized in the following list:
  - Replace instances of `getValue()` with `getAttribute()`
  - Replace instances of `getValueNames()` with `getAttributeNames()`
  - Replace instances of `removeValue()` with `removeAttribute()`
  - Replace instances of `putValue()` with `setAttribute()`
2. During application development, modify Web application behavior as needed.  
In accordance with the Servlet 2.3 specification, HttpSession objects must be scoped within a single Web application context; they may not be shared between contexts. This means that a session can no longer span Web applications. Objects added to a session by a servlet or JSP in one Web application cannot be accessed from another Web application. The same session ID may be shared (because the same cookie is in use), but each Web application will have a unique session associated with the session ID. Version 5 provides a feature that can be used to extend scope of a session to enterprise application.
3. Use administrative tools to configure Session Manager security settings as needed.  
Relative to session security, the default Session Manager setting for Integrate Security is now false. This is different from the default setting in some earlier releases.
4. Use administrative tools to configure the JSP enabler and application server as needed.  
In Version 3.5 of the product, JSP files that contained the usebean tag with scope set to session did not always work properly when session persistence was enabled. Specifically, the JSP writer needed to write a scriplet to explicitly set the attribute (that is, to call `setAttribute()`) if it was changed as part of JSP processing.  
Two new features in Version 5.0 help address this problem:
  - You can set `dosetattribute` to true on the JSP `InitParameter`.
  - You can set the `Write Contents` option to `Write all`.



The differences between the two solutions are summarized in the following table:

|   | Applies to     | Configured at      | Action  |
|---|----------------|--------------------|---|
| dosetattribute set to <i>true</i>             | JSP            | JSP enabler        | Assures that JSP session-scoped beans always call <code>setAttribute()</code> |
| Write Contents option set to <i>Write all</i> | Servlet or JSP | application server | All session data (changed or unchanged) is written to the external location   |

If session persistence is enabled and a class reload for the Web application occurs, the sessions associated with the Web application are maintained in the persistent store and will be available after the reload.

---

## Developing session management in servlets

This information, combined with the coding example `SessionSample.java`, provides a programming model for implementing sessions in your own servlets.

### Steps for this task

1. Get the `HttpSession` object.

To obtain a session, use the `getSession()` method of the `javax.servlet.http.HttpServletRequest` object in the Java Servlet 2.3 API.

When you first obtain the `HttpSession` object, the Session Management facility uses one of three ways to establish tracking of the session: cookies, URL rewriting, or Secure Sockets Layer (SSL) information.

Assume the Session Management facility uses cookies. In such a case, the Session Management facility creates a unique session ID and typically sends it back to the browser as a *cookie*. Each subsequent request from this user (at the same browser) passes the cookie containing the session ID, and the Session Management facility uses this ID to find the user's existing `HttpSession` object.

In Step 1 of the code sample, the `Boolean(create)` is set to `true` so that the `HttpSession` object is created if it does not already exist. (With the Servlet 2.3 API, the `javax.servlet.http.HttpServletRequest.getSession()` method with no boolean defaults to `true` and creates a session if one does not already exist for this user.)

2. Store and retrieve user-defined data in the session.

After a session is established, you can add and retrieve user-defined data to the session. The `HttpSession` object has methods similar to those in `java.util.Dictionary` for adding, retrieving, and removing arbitrary Java objects.

In Step 2 of the code sample, the servlet reads an integer object from the `HttpSession`, increments it, and writes it back. You can use any name to identify values in the `HttpSession` object. The code sample uses the name `sessiontest.counter`.

Because the `HttpSession` object is shared among servlets that the user might access, consider adopting a site-wide naming convention to avoid conflicts.

3. (Optional) Output an HTML response page containing data from the `HttpSession` object.

4. Provide feedback to the user that an action has taken place during the session. You may want to pass HTML code to the client browser indicating that an action has occurred.

For example, in step 3 of the code sample, the servlet generates a Web page that is returned to the user and displays the value of the `sessiontest.counter` each time the user visits that Web page during the session.

5. (Optional) Notify Listeners.

Objects stored in a session that implement the `javax.servlet.http.HttpSessionBindingListener` interface are notified when the session is preparing to end and become invalidated. This notice enables you to perform post-session processing, including permanently saving the data changes made during the session to a database.

6. End the session.

You can end a session:

- Automatically with the Session Management facility if a session is inactive for a specified time. The administrators provide a way to specify the amount of time after which to invalidate a session.
- By coding the servlet to call the `invalidate()` method on the session object.

## SessionSample.java

```
import java.io.*;
import java.util.*;
import javax.servlet.*;
import javax.servlet.http.*;

public class SessionSample extends HttpServlet {
    public void doGet (HttpServletRequest request, HttpServletResponse response)
        throws ServletException, IOException {

        // Step 1: Get the Session object

        boolean create = true;
        HttpSession session = request.getSession(create);

        // Step 2: Get the session data value

        Integer ival = (Integer)
            session.getAttribute ("sessiontest.counter");
        if (ival == null) ival = new Integer (1);
        else ival = new Integer (ival.intValue () + 1);
        session.setAttribute ("sessiontest.counter", ival);

        // Step 3: Output the page

        response.setContentType("text/html");
        PrintWriter out = response.getWriter();
        out.println("<html>");
        out.println("<head><title>Session Tracking Test</title></head>");
        out.println("<body>");
        out.println("<h1>Session Tracking Test</h1>");
        out.println ("You have hit this page " + ival + " times" + "<br>");
        out.println ("Your " + request.getHeader("Cookie"));
        out.println("</body></html>");
    }
}
```

---

## Assembling so that session data can be shared

In accordance with the Servlet 2.3 API specification, by default the Session Management facility supports session scoping by Web module. Only servlets in the same Web module can access the data associated with a particular session. WebSphere Application Server provides an option that you can use to extend the scope of the session attributes to an enterprise application. Therefore, you can share session attributes across all the Web modules in an enterprise application. This option is provided as an IBM extension.

**Restriction:** To use this option, you must install all the Web modules in the enterprise application on a given server. You cannot split up Web modules in the enterprise application by servers. For example, with an enterprise application containing two Web modules, you cannot use this option when one Web module is installed on one server and second Web module is installed on a different server. In such split installations, applications might share session attributes across Web modules using distributed sessions, but session data integrity is lost when concurrent access to a session is made in different Web modules. It also severely restricts use of some Session Management features, like `TIME_BASED_WRITES`. For enterprise applications on which this option is enabled, the Session Management configuration on the Web module inside the enterprise application is ignored. Then Session Management configuration defined on enterprise application is used if Session Management is overwritten at the enterprise application level. Otherwise, the Session Management configuration on the Web container is used.

Do the following to share session data across Web modules in an enterprise application:

### Steps for this task

1. Launch the Application Assembly Tool (AAT).
2. Click the application (EAR file) you want to share.
3. Click the **IBM extension tab**.
4. Click **Shared httpsession context**.
5. Click **Apply**.

Make sure the class definition of attributes put into session are available to all the Web modules in the enterprise application.

6. Save the application (EAR) file.

## Servlet API Behavior

If shared HttpSession context is turned on in an enterprise application, HttpSession listeners defined in all the Web modules inside the enterprise application are invoked for session events. The order of listener invocation is not guaranteed.

---

## Session security support

You can integrate HTTP sessions and security in IBM WebSphere Application Server. When security integration is enabled in the Session Management facility and a session is accessed in a protected resource, you can access that session only in protected resources from then on. You cannot mix secured and unsecured resources accessing sessions when security integration is turned on. Security integration in the Session Management facility is not supported in form-based login with SWAM.

## Security integration rules for HTTP sessions

Only authenticated users can access sessions created in secured pages and are created under the identity of the authenticated user. Only this authenticated user can access these sessions in other secured pages. To protect these sessions from unauthorized users, you cannot access them from an unsecured page.

## Programmatic details and scenarios

IBM WebSphere Application Server maintains the security of individual sessions.

An identity or user name, readable by the `com.ibm.websphere.servlet.session.IBMSession` interface, is associated with a session. An unauthenticated identity is denoted by the user name `anonymous`. IBM WebSphere Application Server includes the `com.ibm.websphere.servlet.session.UnauthorizedSessionRequestException` class, which is used when a session is requested without the necessary credentials.

The Session Management facility uses the WebSphere Application Server security infrastructure to determine the authenticated identity associated with a client HTTP request that either retrieves or creates a session. WebSphere Application Server security determines identity using certificates, LPTA, and other methods.

After obtaining the identity of the current request, the Session Management facility determines whether to return the session requested using a `getSession()` call or not.

The following table lists possible scenarios in which security integration is enabled with outcomes dependent on whether the HTTP request is authenticated and whether a valid session ID and user name was passed to the Session Management facility.

| (Scenario)   | Unauthenticated HTTP request is used to retrieve a session  | HTTP request is authenticated, with an identity of "FRED" used to retrieve a session              |
|--|---|---|
| No session ID was passed in for this request, or the ID is for a session that is no longer valid         | A new session is created. The user name is <code>anonymous</code>                                 | A new session is created. The user name is <code>FRED</code>                                      |
| A session ID for a valid session is passed in. The current session user name is <code>"anonymous"</code> | The session is returned.  | The session is returned. Session Management changes the user name to <code>FRED</code>            |
| A session ID for a valid session is passed in. The current session user name is <code>FRED</code>        | The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown* | The session is returned.  |
| A session ID for a valid session is passed in. The current session user name is <code>BOB</code>         | The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown* | The session is not returned. An <code>UnauthorizedSessionRequestException</code> error is thrown* |

\* A `com.ibm.websphere.servlet.session.UnauthorizedSessionRequestException` error is thrown to the servlet.

---

## Session management support

WebSphere Application Server provides facilities, grouped under the heading *Session Management*, that support the `javax.servlet.http.HttpSession` interface described in the Servlet API specification.

In accordance with the Servlet 2.3 API specification, the Session Management facility supports session scoping by Web module. Only servlets in the same Web module can access the data associated with a particular session. Multiple requests from the same browser, each specifying a unique Web application, result in multiple sessions with a shared session ID. You can invalidate any of the sessions that share a session ID without affecting the other sessions.

You can configure a session timeout for each Web application. A Web application timeout value of 0 (the default value) means that the invalidation timeout value from the Session Management facility is used.

When an HTTP client interacts with a servlet, the state information associated with a series of client requests is represented as an HTTP session and identified by a session ID. Session Management is responsible for managing HTTP sessions, providing storage for session data, allocating session IDs, and tracking the session ID associated with each client request through the use of cookies or URL rewriting techniques. Session Management can store session-related information in several ways:

- In application server memory (the default). This information cannot be shared with other application servers.
- In a database. This storage option is known as *database persistent sessions*.

The last two options are referred to as *distributed sessions*. Distributed sessions are essential for using HTTP sessions for failover facility. When an application server receives a request associated with a session ID that it currently does not have in memory, it can obtain the required session state by accessing the external store (database or memory-to-memory). If distributed session support is not enabled, an application server cannot access session information for HTTP requests that are sent to servers other than the one where the session was originally created. Session Management implements caching optimizations to minimize the overhead of accessing the external store, especially when consecutive requests are routed to the same application server.

Storing session states in an external store also provides a degree of fault tolerance. If an application server goes offline, the state of its current sessions is still available in the external store. This availability enables other application servers to continue processing subsequent client requests associated with that session.

Saving session states to an external location does not completely guarantee their preservation in case of a server failure. For example, if a server fails while it is modifying the state of a session, some information is lost and subsequent processing using that session can be affected. However, this situation represents a very small period of time when there is a risk of losing session information.

The drawback to saving session states in an external store is that accessing the session state in an external location can use valuable system resources. Session Management can improve system performance by caching the session data at the server level. Multiple consecutive requests that are directed to the same server can find the required state data in the cache, reducing the number of times that the

actual session state is accessed in external store and consequently reducing the overhead associated with external location access.

---

## Configuring session management by level

When you configure session management at the Web container level, all applications and the respective Web modules in the Web container normally inherit that configuration, setting up a basic default configuration for the applications and Web modules below it.

However, you can set up different configurations individually for specific applications and Web modules that vary from the Web container default. These different configurations override the default for these applications and Web modules only.

**Note:** When you overwrite the default session management settings on the application level, all the Web modules below that application inherit this new setting unless they too are set to overwrite these settings.

### Steps for this task

1. Open the Administrative console.
2. Select the level that this configuration applies to:
  - For the web container level:
    - a. Click **Servers > Application Servers**.
    - b. Select a server from the list of application servers.
    - c. Under Additional Properties, click **Web Container**.
  - For the enterprise application level:
    - a. Click **Applications > Applications**.
    - b. Select an applications from the list of applications.
  - For the Web module level:
    - a. Click **Applications > Enterprise Applications**.
    - b. Select an applications from the list of applications.
    - c. Under Related Items, click **Web Modules**.
    - d. Select a Web module from the list of Web modules defined for this application.
3. Under **Additional Properties**, click **Session Management**.
4. Make whatever changes you need to manage sessions
5. If you are working on the Web module or application level and want these settings to override the inherited Session Management settings, under **General Properties**, select **Overwrite**.
6. Click **Apply** and **Save**.

---

## Session tracking options

There are several options for session tracking, depending on what sort of tracking method you want to use:

- Session tracking with cookies
- Session tracking with URL rewriting
- Session tracking with SSL information

## Session tracking with cookies

Tracking sessions with cookies is the default. No special programming is required to track sessions with cookies.

## Session tracking with URL rewriting

An application that uses URL rewriting to track sessions must adhere to certain programming guidelines. The application developer needs to do the following:

- Program servlets to encode URLs
- Supply a servlet or Java Server Pages (JSP) file as an entry point to the application

Using URL rewriting also requires that you enable URL rewriting in the Session Management facility.

**Note:** In certain cases, clients cannot accept cookies. Therefore, you cannot use cookies as a session tracking mechanism. Applications can use URL rewriting as a substitute.

### Program session servlets to encode URLs

Depending on whether the servlet is returning URLs to the browser or redirecting them, include either `encodeURL()` or `encodeRedirectURL()` in the servlet code. Examples demonstrating what to replace in your current servlet code follow.

#### Rewrite URLs to return to the browser

Suppose you currently have this statement:

```
out.println("<a href=\"/store/catalog\">catalog<a>");
```

Change the servlet to call the `encodeURL` method before sending the URL to the output stream:

```
out.println("<a href=\"\"");  
out.println(response.encodeURL ("/store/catalog"));  
out.println(">catalog</a>");
```

#### Rewrite URLs to redirect

Suppose you currently have the following statement:

```
response.sendRedirect ("http://myhost/store/catalog");
```

Change the servlet to call the `encodeRedirectURL` method before sending the URL to the output stream:

```
response.sendRedirect (response.encodeRedirectURL ("http://myhost/store/catalog"));
```

The `encodeURL()` and `encodeRedirectURL()` methods are part of the `HttpServletResponse` object. These calls check to see if URL rewriting is configured before encoding the URL. If it is not configured, the calls return the original URL.

If both cookies and URL rewriting are enabled and `response.encodeURL()` or `encodeRedirectURL()` is called, the URL is encoded, even if the browser making the HTTP request processed the session cookie.

You can also configure session support to enable protocol switch rewriting. When this option is enabled, the product encodes the URL with the session ID for switching between HTTP and HTTPS protocols.

### Supply a servlet or JSP file as an entry point

The entry point to an application (such as the initial screen presented) may not require the use of sessions. However, if the application in general requires session support (meaning some part of it, such as a servlet, requires session support), then after a session is created, all URLs are encoded to perpetuate the session ID for the servlet (or other application component) requiring the session support.

The following example shows how you can embed Java code within a JSP file:

```
<%  
response.encodeURL ("/store/catalog");  
%>
```

## Session tracking with SSL information

No special programming is required to track sessions with Secure Sockets Layer (SSL) information.

To use SSL information, turn on **Enable SSL ID tracking** in the Session Management property page. Because the SSL session ID is negotiated between the Web browser and HTTP server, this ID cannot survive an HTTP server failure. However, the failure of an application server does not affect the SSL session ID if an external HTTP Server is present between WebSphere Application Server and the browser.

SSL tracking is supported for the IBM HTTP Server and iPlanet Web servers only. You can control the lifetime of an SSL session ID by configuring options in the Web server. For example, in the IBM HTTP Server, set the configuration variable SSLV3TIMEOUT to provide an adequate lifetime for the SSL session ID. An interval that is too short can cause a premature termination of a session. Also, some Web browsers might have their own timers that affect the lifetime of the SSL session ID. These Web browsers may not leave the SSL session ID active long enough to serve as a useful mechanism for session tracking. Internal Http Server of WebSphere also supports SSL Tracking.

When using the SSL session ID as the session tracking mechanism in a cloned environment, use either cookies or URL rewriting to maintain session affinity. The cookie or rewritten URL contains session affinity information that enables the Web server to properly route a session back to the same server for each request.

---

## Configuring session tracking

To configure session tracking, complete the following:

### Steps for this task

1. Go to the appropriate level of Session Management.
2. Specify which session tracking mechanism you want to pass the session ID between the browser and the servlet:
  - To track sessions with cookies, click **Enable Cookies**.  
To change the cookie settings, click **Modify**.
  - To track sessions with URL rewriting, click **Enable URL Rewriting**.



If you want to enable protocol switch rewriting, click **Enable protocol switch rewriting**.

- To track sessions with SSL information, click **Enable SSL ID tracking**.

3. Click **Apply**.
4. Click **Save**.
5. Define the session recovery characteristics.

## Serializing access to session data

The Servlet API supports concurrent access to a session in a given server instance. WebSphere Application Server provides an option to prevent the concurrent access to a session in a given server instance so that concurrent modification of a session does not occur in a given server instance. This prevention is achieved by synchronizing the requests based on session. When this feature is turned on, a session is obtained for the request before invoking the servlet and requests are synchronized by locking the session for the servlet execution time. Note that synchronization is based on the memory copy of session. So this feature cannot serialize requests across servers based on session when session affinity fails.

**Restriction:** Use this feature only when concurrent modification of the same session data is possible and is not desirable by the application. This feature has overhead of serializing the requests based on a session.

Do the following to synchronize session access:

### Steps for this task

1. Select the level of Session Management on which you want to serialize session access.
2. Under Serialize Session access, click **Allow serial access**.
3. In the Maximum wait time box, type the amount of time, in milliseconds, a servlet waits on a session before continuing execution. The default is 120000 milliseconds or two minutes.
4. **(Optional)** Select **Allow access on timeout** if you want the servlet to gain access to the session and continue normal execution even if the session is still locked by another servlet. If you do not select this box, the servlet execution will abort when the session request times out.
5. Click **Apply**.
6. Click **Save**.

## Session Management settings

Use this page to manage HTTP session support. This support includes specifying a session tracking mechanism, setting maximum in-memory session count, controlling overflow, and configuring session timeout.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Web Container > Session Management**.

### **Overwrite Session Management (application and Web module level only)**

Specifies whether or not these Session Management settings take precedence over those normally inherited from a higher level for the current application or Web module.

By default, Web modules inherit Session Management settings from the application level above it, and applications inherit Session Management settings from the Web container level above it.

### Session tracking mechanism

Specifies a mechanism for HTTP session management.

| Mechanism              | Function  | Default      |
|------------------------|---|--------------|
| Enable SSL ID Tracking | Specifies that session tracking uses Secure Sockets Layer (SSL) information as a session ID. Enabling SSL tracking takes precedence over cookie-based session tracking and URL rewriting. | 9600 seconds |

There are two parameters available if you enable SSL ID tracking: `SSLV3Timeout` and `Secure Association Server (SAS)`. `SSLV3Timeout` specifies the time interval after which SSL sessions are renegotiated. This is a high setting and modification does not provide any significant impact on performance. The `SAS` parameter establishes an SSL connection only if it goes out of the Java Virtual Machine (JVM) to another JVM. If all the beans are co-located within the same JVM, the SSL used by SAS does not hinder performance.

These are set by editing the `sas.server.properties` and `sas.client.props` files located in the `product_installation_root\properties` directory, where `product_installation_root` is the directory where WebSphere Application Server is installed.

Enable Cookies

Specifies that session tracking uses cookies to carry session IDs. If cookies are enabled, session tracking recognizes session IDs that arrive as cookies and tries to use cookies for sending session IDs. If cookies are not enabled, session tracking uses URL rewriting instead of cookies (if URL rewriting is enabled).

Enabling cookies takes precedence over URL rewriting. Do not disable cookies in the Session Management facility of the application server that is running the administrative application because this action causes the administrative application not to function after a restart of the server. As an alternative, run the administrative application in a separate process from your applications.

Enable URL Rewriting

Click **Modify** to change these settings.

Specifies that the Session Management facility uses rewritten URLs to carry the session IDs. If URL rewriting is enabled, the Session Management facility recognizes session IDs that arrive in the URL if the `encodeURL` method is called in the servlet.

Enable Protocol Switch Rewriting

Specifies that the session ID is added to a URL when the URL requires a switch from HTTP to HTTPS or from HTTPS to HTTP. If rewriting is enabled, the session ID is required to go between HTTP and HTTPS.

### **Overflow**

Specifies that the number of sessions in memory can exceed the value specified by the Max In Memory Session Count property. This option is valid only in nondistributed sessions mode.

### **Maximum in-memory session count**

Specifies the maximum number of sessions to maintain in memory.

The meaning differs depending on whether you are using in-memory or distributed sessions. For in-memory sessions, this value specifies the number of sessions in the base session table. Use the Allow Overflow property to specify whether to limit sessions to this number for the entire Session Management facility or to allow additional sessions to be stored in secondary tables. For distributed sessions, this value specifies the size of the memory cache for sessions. When the session cache has reached its maximum size and a new session is requested, the Session Management facility removes the least recently used session from the cache to make room for the new one.

## 5.0.2

**Note:** Do not set this value to a number less than the maximum thread pool size for your server.

### Session timeout

Specifies how long a session can go unused before it is no longer valid. Specify either Set timeout or No timeout. Specify the value in minutes greater than or equal to two.

The value of this setting is used as a default when the session timeout is not specified in a Web module deployment descriptor. Note that to preserve performance, the invalidation timer is not accurate to the second. When the Write Frequency is time based, ensure that this value is least twice as large as the write interval.

### Security integration

Specifies that when security integration is enabled, the Session Management facility associates the identity of users with their HTTP sessions

### Serialize session access

Specifies that concurrent session access in a given server is not allowed.

Maximum wait time

Specifies the maximum amount of time a servlet request waits on an HTTP session before continuing execution. This parameter is optional and expressed in seconds. The default is 120, or 2 minutes. Under normal conditions, a servlet request waiting for access to an HTTP session gets notified by the request that currently owns the given HTTP session when the request finishes.

Allow access on timeout

Specifies whether the servlet is executed normally or aborted in the event of a timeout. If this box is checked, the servlet executes normally. If this box is not checked, the servlet execution aborts and error logs are generated.

## Cookie settings

Use this page to configure cookie settings for session management.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Web Container > Session Management > Enable Cookies.**

### **Cookie name**

Specifies a unique name for the session management cookie. The servlet specification requires the name JSESSIONID. However, for flexibility this value can be configured.

### **Secure cookies**

Specifies that the session cookies include the secure field. Enabling the feature restricts the exchange of cookies to HTTPS sessions only.

### **Cookie domain**

Specifies the domain field of a session tracking cookie. This value controls whether or not a browser sends a cookie to particular servers. For example, if you specify a particular domain, session cookies are sent to hosts in that domain. The default domain is the server.

### **Cookie path**

Specifies that a cookie is sent to the URL designated in the path. Specify any string representing a path on the server. "/" indicates root directory. Specify a value to restrict the paths to which the cookie will be sent. By restricting paths, you prevent the cookie from going to certain URLs on the server. If you specify the root directory, the cookie is sent no matter which path on the given server is accessed.

### **Cookie maximum age**

Specifies the amount of time that the cookie lives on the client browser. Specify that the cookie lives only as long as the current browser session, or to a maximum age. If you choose the maximum age option, specify the age in seconds. This value corresponds to the Time to Live (TTL) value described in the Cookie specification.

Default is the current browser session which is equivalent to setting the value to -1.

---

## **Configuring session tracking for Wireless Application Protocol (WAP) devices**

Most Wireless Application Protocol (WAP) devices do not support cookies. The preferred way to track sessions for WAP devices is to use URL rewriting. However on most WAP devices, the maximum allowed URL length is 128 characters. With URL rewriting, a session identifier is added to the URL itself, effectively decreasing the space available for the actual URL and the number of parameters that can be sent on a request.

To reduce the length of session identifier, you can configure key (jsessionid), session ID length and clone ID. To make these configuration changes, complete the following:

#### Steps for this task

1. Open the Administrative console.
2. Click **Servers > Application Servers**.
3. Select a server from the list of application servers.
4. Under Additional Properties, click **Web Container**
5. Under Additional Properties, click **Custom Properties**.
6. Add the appropriate properties from the following list:
  - HttpSessionIdLength
  - SessionRewriteIdentifier

- HttpSessionCloneId
- CloneSeparatorChange
- NoAdditionalSessionInfo
- SessionIdentifierMaxLength

7. Click **Apply** and **Save**.

## Session management custom properties

Custom properties for session management:

### CloneSeparatorChange

Use this property to maintain session affinity. The clone ID of the server is appended to session identifier separated by colon. On some Wireless Application Protocol (WAP) devices, a colon is not allowed. Set this property to "true" to change clone separator to a plus sign (+).

### HttpSessionCloneId

Use this property to change the clone ID of the cluster member. Within a cluster, this name must be unique to maintain session affinity. When set, this name overwrites the default name generated by WebSphere Application Server. Default clone ID length: 8 or 9.

### HttpSessionIdLength

Use this property to configure the session identifier length. Do not use an extremely low value; using a low value results in reduced number of combinations possible, thereby increasing risk of guessing the session identifier. In a cluster, all cluster members should be configured with same ID length. Allowed range: 8 to 128. Default length: 23.

### HttpSessionReaperPollInterval

Use this property to set a wake-up interval for the process that removes invalid sessions. Default is based on maximum inactive interval set in Session Management. Allowed value: integer.

### NoAdditionalSessionInfo

Set this value to "true" to force removal of information that is not needed in session identifiers. In WebSphere Application Server base edition, a clone ID of -1 is never used; therefore, a clone ID is not included in base edition when this is set. Also, cache ID is not used with nonpersistent sessions; so the cache ID is not included with nonpersistent sessions when this value is set.

### SessionIdentifierMaxLength

Use this value to set maximum length that a session identifier can grow. In a cluster, because of fail-over when a request goes to new cluster member, Session Management appends a new clone ID to the existing clone ID. In a large cluster, if for some reason servers are failing more often, then it is possible that the session identifier length can be more than expected reducing room for URL. So this property helps to find out the condition and take appropriate action to address servers fail-over. When this is specified, message is logged when specified maximum length is reached. Allowed value: integer.

### SessionRewriteIdentifier

Use this property to change the key used with URL rewriting. Default key: jsessionid.

---

## Distributed sessions

WebSphere Application Server provides the following session mechanisms in a distributed environment:

- **Database Session persistence**, where sessions are stored in the database specified.

When a session contains attributes that implement `HttpSessionActivationListener`, notification occurs anytime the session is activated (that is, session is read to the memory cache) or passivated (that is, session leaves the memory cache). Passivation can occur because of a server shutdown or when the session memory cache is full and an older session is removed from the memory cache to make room for a newer session. It is not guaranteed that a session is passivated in one application server prior to being activated in another.

---

## Session recovery support

For session recovery support, WebSphere Application Server provides distributed session support in the form of database sessions. Use session recovery support under the following conditions:

- When the user's session data must be maintained across a server restart
- When the user's session data is too valuable to lose through an unexpected server failure

All the attributes set in a session must implement `java.io.Serializable` if the session requires external storage. In general, consider making all objects held by a session serialized, even if immediate plans do not call for session recovery support. If the Web site grows, and session recovery support becomes necessary, the transition occurs transparently to the application if the sessions only hold serialized objects. If not, a switch to session recovery support requires coding changes to make the session contents serialized.

## Distributed Environment settings

Use this page to specify a type for saving a session in a distributed environment.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Web Container > Session Management > Distributed Environment Settings**.

### Distributed Sessions

Specifies the type of distributed environment to be used for saving sessions.

None

Specifies that the Session Management facility discards the session data when the server shuts down.

Database

Specifies that the Session Management facility stores session information in the data source specified by the data source connection settings. Click **Database** to change these data source settings.

---

## Configuring for database session persistence

To configure the Session Management facility for database session persistence, complete the following:

### Steps for this task

1. Define a JDBC provider.
2. Create a data source pointing to an existing database, using the JDBC provider that you defined. Note the JNDI name of the data source.  

Under **Data Sources** > *datasource\_name* > **Custom Properties**, make sure the correct DB2 database is entered for the value of the **databaseName** property. If necessary, contact your DB2 Database Administrator to verify the correct database name.
3. Go to the appropriate level of Session Management.
4. Click **Distributed Environment Settings**
5. Select and click **Database**.
6. Specify the Data Source JNDI name from step 2.
7. Specify the database user ID and password for accessing the database.
8. Retype the password for confirmation.
9. **(Optional)** Configure a table space and page sizes for DB2 session databases.
10. **(Optional)** Switch to a multirow schema.
11. Click **OK**.
12. **(Optional)** If you want to change the tuning parameters, click **Custom Tuning Parameters** and select a setting or customize one.
13. Click **Apply**.
14. Click **Save**.

## Switching to a multirow schema

By default, a single session maps to a single row in the database table used to hold sessions. With this setup, there are hard limits to the amount of user-defined, application-specific data that WebSphere Application Server can access.

### Steps for this task

1. Modify the Session Management facility properties to switch from single to multirow schema.
2. Manually drop the database table or delete all the rows in the database table that the product uses to maintain HttpSession objects.  

To drop the table:

  - a. Determine which data source configuration Session Management is using.
  - b. In the data source configuration, look up the database name.
  - c. Use the database facilities to connect to the database.
  - d. Drop the SESSIONS table.

## Configuring tablespace and page sizes for DB2 session databases

If you are using DB2 for session persistence, you can increase the page size to optimize performance for writing large amounts of data to the database. Page sizes of 8K, 16K, and 32K are supported.



To use a page size other than the default (4K), do the following:

#### Steps for this task

1. If the SESSIONS table already exists, drop it from the DB2 database.
2. Create a new DB2 buffer pool and table space, specifying the same page size (8K, 16K or 32K) for both, and assign the new buffer pool to this table space.

```
DB2 Connect to session
DB2 CREATE BUFFERPOOL sessionBP SIZE 1000 PAGESIZE 8K
DB2 Connect reset
DB2 Connect to session
DB2 CREATE TABLESPACE sessionTS PAGESIZE 8K MANAGED BY SYSTEM
    USING ('D:\DB2\NODE0000\SQL00005\sessionTS.0') BUFFERPOOL sessionBP
DB2 Connect reset
```

Refer to DB2 product documentation for details.

3. Configure the correct table space name and page size in the Session Management facility.

Page size is referred to as *row size* on the Session Management page.)

#### Results

When the product is restarted, the Session Management facility creates the new SESSIONS table in the specified tablespace based on the indicated page size.

## Database settings

Use this page to specify the settings for database session support.

To view this administrative console page, click **Servers > Application Servers > server\_name > Web Container > Session Management > Distributed Environment Settings > Database**.

### **Datasource JNDI Name**

Specifies the datasource description

The JNDI name of the non-XA enabled data source from which Session Management obtains database connections. For example, if the JNDI name of the datasource is "jdbc/sessions", specify "jdbc/sessions." The data source represents a pool of database connections and a configuration for that pool (such as the pool size). The data source must already exist as a configured resource in the environment.

### **User ID**

Specifies the user ID for database access

### **Password**

Specifies the password for database access

### **Confirm Password**

Specifies the password a second time to ensure it recorded correctly.

### **DB2 Row Size**

Specifies the tablespace page size configured for the sessions table, if using a DB2 database. Possible values are 4, 8, 16, and 32 kilobytes (K). The default row size is 4K.

The default row size is 4K. In DB2, it can be updated to a larger value. This can help database performance in some environments. When this value is other than 4, you must specify Table Space Name to use. For 4K pages, the Table Space Name is optional.

### Table Space Name

Specifies that tablespace to be used for the sessions table.

This value is required when the DB2 Page Size is other than 4K.

### Use Multirow Sessions

Specifies that each instance of application data be placed in a separate row in the database, allowing larger amounts of data to be stored for each session. This action can yield better performance in certain usage scenarios. If using multirow schema is not enabled, instances of application data can be placed in the same row.

## Multirow schema considerations

IBM WebSphere Application Server supports the use of a multirow schema option in which each piece of application specific data is stored in a separate row of the database. With this setup, the total amount of data you can place in a session is now bound only by the database capacities. The only practical limit that remains is the size of the session attribute object.

The multirow schema potentially has performance benefits in certain usage scenarios, such as when larger amounts of data are stored in the session but only small amounts are specifically accessed during a given servlet processing of an HTTP request. In such a scenario, avoiding unneeded Java object serialization is beneficial to performance.

Understand that switching between multirow and single row is not a trivial proposition.

In addition to allowing larger session records, using multirow schema can yield performance benefits. However, it requires a little work to switch from single-row to multirow schema, as shown in the instructions below.

### Coding considerations and test environment

Consider configuring direct single-row usage to one database and multirow usage to another database while you verify which option suits your application needs. (Do this in code by switching the data source used; then monitor performance.)

| Programming issue                    | Application scenario   |
|--------------------------------------|--|
| Reasons to use single-row            | <ul style="list-style-type: none"><li>You can read or write all values with just one record read and write.</li><li>This takes up less space in a database because you are guaranteed that each session is only one record long.</li></ul> |
| Reasons <b>not</b> to use single-row | 2-megabyte limit of stored data per session.   |

| Programming issue                  | Application scenario   |
|------------------------------------|--|
| Reasons to use multirow            | <ul style="list-style-type: none"> <li>• The application can store an unlimited amount of data; that is, you are limited only by the size of the database and a 2-megabyte-per-record limit.</li> <li>• The application can read individual fields instead of the whole record. When large amounts of data are stored in the session but only small amounts are specifically accessed during servlet processing of an HTTP request, multirow sessions can improve performance by avoiding unneeded Java object serialization.</li> </ul> |
| Reasons <b>not</b> to use multirow | If data is small in size, you probably do not want the extra overhead of multiple row reads when you can store everything in one row.  |

In the case of multirow usage, design your application data objects not to have references to each other, to prevent circular references. For example, suppose you are storing two objects A and B in the session using `HttpSession.put(..)` method, and A contains a reference to B. In the multirow case, because objects are stored in different rows of the database, when objects A and B are retrieved later, the object graph between A and B is different than stored. A and B behave as independent objects.

---

## Clustered session support

A clustered environment supports load balancing, where the workload is distributed among the application servers that compose the cluster. In a cluster environment, the same Web application must exist on each of the servers that can access the session. You can accomplish this setup by installing an application onto a cluster definition. Each of the servers in the group can then access the Web application

In a clustered environment, the Session Management facility requires an affinity mechanism so that all requests for a particular session are directed to the same application server instance in the cluster. This requirement conforms to the Servlet 2.3 specification in that multiple requests for a session cannot coexist in multiple application servers. One such solution provided by IBM WebSphere Application Server is *session affinity* in a cluster; this solution is available as part of the WebSphere Application Server plug-ins for Web servers. It also provides for better performance because the sessions are cached in memory. In clustered environments other than WebSphere Application Server clusters, you must use an affinity mechanism (for example, IBM WebSphere Edge Server affinity).

If one of the servers in the cluster fails, it is possible for the request to reroute to another server in the cluster. If distributed sessions support is enabled, the new server can access session data from the database or another WebSphere Application Server instance. You can retrieve the session data only if a new server has access to an external location from which it can retrieve the session.

---

## Tuning session management

IBM WebSphere Application Server session support has features for tuning session performance and operating characteristics, particularly when sessions are configured in a distributed environment. These options support the administrator flexibility in determining the performance and failover characteristics for their environment.

The table summarizes the features, including whether they apply to sessions tracked in memory, in a database, with memory-to-memory replication, or all. Click a feature for details about the feature. Some features are easily manipulated using administrative settings; others require code or database changes.

| Feature or option                | Goal  | Applies to sessions in memory, database, or memory-to-memory |
|----------------------------------|---|--|
| Write frequency                  | Minimize database write operations.   | Database and Memory-to-Memory                                |
| Session affinity                 | Access the session in the same application server instance.   | All  |
| Multirow schema                  | Fully utilize database capacities.  | Database   |
| Base in-memory session pool size | Fully utilize system capacity without overburdening system.   | All  |
| Write contents                   | Allow flexibility in determining what session data to write   | Database and Memory-to-Memory                                |
| Scheduled invalidation           | Minimize contention between session requests and invalidation of sessions by the Session Management facility. Minimize write operations to database for updates to last access time only. | Database and Memory-to-Memory                                |
| Tablespace and row size          | Increase efficiency of write operations to database.  | Database (DB2 only)  |

### Configuring scheduled invalidation

You can set specific times for the Session Management facility to scan for invalidated sessions in a distributed environment. When used with distributed sessions, this feature has the following benefits:

- You can schedule the scan for invalidated sessions for times of low application server activity, avoiding contention between invalidation scans of database or another WebSphere Application Server instance and read and write operations to service HTTP session requests.
- Significantly fewer external write operations can occur when running with the End of Service Method write mode because the last access time of the session does not need to be written out on each HTTP request. (Manual Update and Time Based Write options already minimize the writing of the last access time.)

## Usage considerations

- With scheduled invalidation configured, HttpSession timeouts are not strictly enforced. Instead, all invalidation processing is handled at the configured invalidation times.
- HttpSessionBindingListener processing is handled at the configured invalidation times unless the HttpSession.invalidate( ) method is explicitly called.
- The HttpSession.invalidate() method immediately invalidates the session from both the session cache and the external store.

## Configuring write contents

In Session Management, you can configure which session data is written to the database or to another WebSphere instance, depending on whether you are using database persistent sessions or memory to memory replication. This flexibility allows for fewer code changes for the JSP writer when the application will be operating in a clustered environment. The following options are available in Session Management for tuning what is to be written back:

- Write changed (the default) - Write only session data properties that have been updated through setAttribute() and removeAttribute() method calls.
- Write all - Write all session data properties.

The **Write all** setting might benefit servlet and JSP writers who change Java objects' states that reside as attributes in HttpSession and do not call HttpSession.setAttribute().

However, the use of **Write all** could result in more data being written back than is necessary. If this situation applies to you, consider combining the use of **Write all** with **Time-based write** to boost performance overall. As always, be sure to evaluate the advantages and disadvantages for your installation.

With either Write Contents setting, when a session is first created, complete session information is written, including all of the objects bound to the session. When using database session persistence, in subsequent session requests, what is written to the database depends on whether a single-row or multirow schema has been set for the session database, as follows:

| Write Contents setting | Behavior with single-row schema  | Behavior with multirow schema   |
|------------------------|--|---|
| Write changed          | If any session attribute is updated, all objects bound to the session are written. | Only the session data modified through setAttribute() or removeAttribute() calls is written.  |
| Write all              | All bound session attributes are written.  | All session attributes that currently reside in the cache are written. If the session has never left the cache, all session attributes are written. |

### Steps for this task

1. Go to the appropriate level of Session Management.
2. Click Distributed Environment Settings
3. Click Custom Tuning Parameters.
4. Select Custom Settings, and click Modify.

5. Select the appropriate write contents setting.

## Configuring write frequency

In the Session Management facility, you can configure the frequency for writing session data to the database or to a WebSphere instance, depending on whether you use database distributed sessions or memory-to-memory replication. This flexibility enables you to weigh session performance gains against varying degrees of failover support. The following options are available in the Session Management facility for tuning write frequency:

- `END_OF_SERVICE` - Write session data at the end of the servlet `service()` method call.
- `MANUAL_UPDATE` - Write session data only when the servlet calls the `IBMSession.sync()` method.
- `TIME_BASED_WRITE` (the default) - Write session data at periodic intervals, in seconds (called the *write interval*).

When a session is first created, session information is always written at the end of the `service()` call.

## Base in-memory session pool size

The base in-memory session pool size number has different meanings, depending on session support configuration:

- With in-memory sessions, session access is optimized for up to this number of sessions.
- With distributed sessions (meaning, when sessions are stored in a database or in another WebSphere Application Server instance); it also specifies the cache size and the number of last access time updates saved in manual update mode.

For distributed sessions, when the session cache has reached its maximum size and a new session is requested, the Session Management facility removes the least recently used session from the cache to make room for the new one.

General memory requirements for the hardware system, and the usage characteristics of the e-business site, determines the optimum value.

Note that increasing the base in-memory session pool size can necessitate increasing the heap sizes of the Java processes for the corresponding WebSphere Application Servers.

### Overflow in nondistributed sessions

By default, the number of sessions maintained in memory is specified by base in-memory session pool size. If you do not wish to place a limit on the number of sessions maintained in memory and allow overflow, set `overflow` to *true*.

Allowing an unlimited amount of sessions can potentially exhaust system memory and even allow for system sabotage. Someone could write a malicious program that continually hits your site and creates sessions, but ignores any cookies or encoded URLs and never utilizes the same session from one HTTP request to the next.

When overflow is disallowed, the Session Management facility still returns a session with the `HttpServletRequest.getSession(true)` method when the memory limit is reached, and this is an invalid session that is not saved.

With the WebSphere Application Server extension to HttpSession, `com.ibm.websphere.servlet.session.IBMSession`, an `isOverflow()` method returns *true* if the session is such an invalid session. An application can check this status and react accordingly.

## Controlling write operations

You can manually control when modified session data is written out to the database or to another WebSphere Application Server instance by using the `sync()` method in the `com.ibm.websphere.servlet.session.IBMSession` interface, which extends the `javax.servlet.http.HttpSession` interface. By calling the `sync()` method from the `service()` method of a servlet, you send any changes in the session to the external location. When `MANUAL_UPDATE` is selected as the write frequency mode, session data changes are written to an external location only if the application calls the `sync()` method. If the `sync()` method is not called, session data changes are lost when a session object leaves the server cache. With `END_OF_SERVICE` or `TIME_BASE_WRITE` is the write frequency mode, the session data changes are written out whenever the `sync()` method is called. If the `sync()` method is not called, changes are written out at the end of service method or on a time interval basis based on the write frequency mode selected.

```
IBMSession iSession = (IBMSession) request.getSession();
iSession.setAttribute("name", "Bob");

//force write to external store
iSession.sync( )
```

## Tuning parameter settings

Use this page to set tuning parameters for distributed sessions.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Web Container > Session Management > Distributed Environment Settings > Custom Tuning Parameters**.

### Tuning Level

Specifies that the Session Management facility provides certain predefined settings that affect performance.

Select one of these predefined settings or customize a setting. The default setting is 10 seconds.

To customize a setting, select one of the predefined settings that comes closest to the setting you want, click **Custom settings**, make your changes, and then click **OK**.

Predefined options and their settings follow:

#### Very high (optimize for performance):

|                            |                         |
|----------------------------|-------------------------|
| Write frequency            | TIME_BASED_WRITES       |
| Write interval             | 300 sec                 |
| Write Contents             | Only updated attributes |
| Schedule Sessions Clean up | true                    |
| First Hour                 | 00                      |

#### High:

|                 |                        |
|-----------------|------------------------|
| Write frequency | TIME_BASED_WRITES      |
| Write interval  | 300 sec                |
| Write Contents  | All session attributes |

**Medium:**

|                 |                         |
|-----------------|-------------------------|
| Write frequency | END_OF_SERVICE          |
| Write Contents  | Only updated attributes |

**Low (optimize for failover)**

|                 |                        |
|-----------------|------------------------|
| Write frequency | END_OF_SERVICE         |
| Write Contents  | All session attributes |

**Write frequency**

Specifies when the session writes to the database or another WebSphere Application Server instance happens.

|                        |  |
|------------------------|--|
| End of servlet service | A session writes to a database or another WebSphere Application Server instance after the servlet completes execution.                               |
| Manual update          | A programmatic sync on the IBMSession object is required to write the session data to the database or another WebSphere Application Server instance. |
| Time based             | Session data writes to the database or another WebSphere Application Server instance based on the specified Write Interval value.                    |

**Write contents**

Specifies whether updated attributes are only written to the external location or all of the session attributes are written to the external location, regardless of whether or not they changed. The external location can be either a database or another application server instance.

**Schedule sessions cleanup**

Specifies when to clean the invalid sessions from a database or another application server instance.

**Specify sessions cleanup schedule**

Enables the scheduled invalidation process for cleaning up the invalidated HTTP sessions from the external location. Enable this option to reduce the number of updates to a database or another application server instance required to keep the HTTP sessions alive. When this option is not enabled, the invalidator process runs every few minutes to remove invalidated HTTP sessions.

When this option is enabled, specify the two hours of a day for the process to clean up the invalidated sessions in the external location. Specify the times when there is the least activity in the application servers. An external location can be either a database or another application server instance.

**First Time of Day**



Indicates the first hour, in Greenwich Mean Time (GMT), during which the invalidated sessions are cleared from the external location. Specify this value as a positive integer between 0 and 23. This value is valid only when schedule invalidation is enabled.

#### Second Time of Day

Indicates the second hour, in Greenwich Mean Time (GMT), during which the invalidated sessions are cleared from the external location. Specify this value as a positive integer between 0 and 23. This value is valid only when schedule invalidation is enabled.

---

## Best practices for using HTTP Sessions

- **Enable Security integration for securing HTTP sessions**

HTTP sessions are identified by session IDs. A session ID is a pseudo-random number generated at the runtime. Session hijacking is a known attack HTTP sessions and can be prevented if all the requests going over the network are enforced to be over a secure connection (meaning, HTTPS). But not every configuration in a customer environment enforces this constraint because of the performance impact of SSL connections. Due to this relaxed mode, HTTP session is vulnerable to hijacking and because of this vulnerability, WebSphere Application Server has the option to tightly integrate HTTP sessions and WebSphere Application Server security. Enable security in WebSphere Application Server so that, the sessions are protected in a manner that only users who created the sessions are allowed to access them.

- **Release HttpSession objects using `javax.servlet.http.HttpSession.invalidate()` when finished.**

HttpSession objects live inside the Web container until:

- The application explicitly and programmatically releases it using the `javax.servlet.http.HttpSession.invalidate()` method; quite often, programmatic invalidation is part of an application logout function.
- WebSphere Application Server destroys the allocated HttpSession when it expires (default = 1800 seconds or 30 minutes). The WebSphere Application Server can only maintain a certain number of HTTP sessions in memory based on Session Management settings. In case of distributed sessions, when maximum cache limit is reached in memory, the Session Management facility removes the least recently used (LRU) one from cache to make room for a session.

- **Avoid trying to save and reuse the HttpSession object outside of each servlet or JSP file.**

The HttpSession object is a function of the HttpRequest (you can get it only through the `req.getSession()` method), and a copy of it is valid only for the life of the `service()` method of the servlet or JSP file. You *cannot* cache the HttpSession object and refer to it outside the scope of a servlet or JSP file.

- **Implement the Serializable class when developing new objects to be stored in the HTTP session.**

This action allows the object to properly serialize when using distributed sessions. Without this extension, the object cannot serialize correctly and throws an error. An example of this follows:

```
public class MyObject implements java.io.Serializable {...}
```

Make sure all instance variable objects that are not marked transient are serializable.

- **The HttpSession API does not dictate transactional behavior for sessions.**

Distributed HttpSession support does not guarantee transactional integrity of an attribute in a failover scenario or when session affinity is broken. Use transactionally aware resources like enterprise Java beans to guarantee the transaction integrity required by your application.

- **Ensure the Java objects you add to a session are in the correct class path.**

If you add Java objects to a session, place the class files for those objects in the correct classpath (the Application Classpath if utilizing sharing across Web modules in an enterprise application, or the WebModule Classpath if using the Servlet 2.2-complaint session sharing) or in the directory containing other servlets used in WebSphere Application Server. In the case of session clustering, this action applies to every node in the cluster.

Because the HttpSession object is shared among servlets that the user might access, consider adopting a site-wide naming convention to avoid conflicts.

- **Avoid storing large object graphs in the HttpSession object.**

In most applications each servlet only requires a fraction of the total session data. However, by storing the data in the HttpSession object as one large object, an application forces WebSphere Application Server to process all of it each time.

- **Utilize Session Affinity to help achieve higher cache hits in the WebSphere Application Server.**

WebSphere Application Server has functionality in the HTTP Server plug-in to help with session affinity. The plug-in will read the cookie data (or encoded URL) from the browser and helps direct the request to the appropriate application or clone based on the assigned session key. This functionality increases use of the in-memory cache and reduces hits to the database or another WebSphere Application Server instance

- **Maximize use of session affinity and avoid breaking affinity.**

Using session affinity properly can enhance the performance of the WebSphere Application Server. Session affinity in the WebSphere Application Server environment is a way to maximize the in-memory cache of session objects and reduce the amount of reads to the database or another WebSphere Application Server instance. Session affinity works by caching the session objects in the server instance of the application with which a user is interacting. If the application is deployed in multiple servers of a server group, the application can direct the user to any one of the servers. If the users starts on server1 and then comes in on server2 a little later, the server must write all of the session information to the external location so that the server instance in which server2 is running can read the database. You can avoid this database read using session affinity. With session affinity, the user starts on server1 for the first request; then for every successive request, the user is directed back to server1. Server1 has to look only at the cache to get the session information; server1 never has to make a call to the session database to get the information.

You can improve performance by not breaking session affinity. Some suggestions to help avoid breaking session affinity are:

- Combine all Web applications into a single application server instance, if possible, and use modeling or cloning to provide failover support.
- Create the session for the frame page, but do not create sessions for the pages within the frame when using multiframe JSP files. (See discussion later in this topic.)

- **When using multi-framed pages, follow these guidelines:**

- Create a session in only one frame or before accessing any frame sets. For example, assuming there is no session already associated with the browser and a user accesses a multi-framed JSP file, the browser issues concurrent requests for the JSP files. Because the requests are not part of any session, the JSP files end up creating multiple sessions and all of the cookies are sent back to the browser. The browser honors only the last cookie that arrives. Therefore, only the client can retrieve the session associated with the last cookie. Creating a session before accessing multi-framed pages that utilize JSP files is recommended.
- By default, JSPs get a HttpSession using `request.getSession(true)` method. So by default JSPs create a new session if none exists for the client. Each JSP page in the browser is requesting a new session, but only one session is used per browser instance. A developer can use `<% @ page session="false" %>` to turn off the automatic session creation from the JSP files that will not access the session. Then if the page needs access to the session information, the developer can use `<%HttpSession session = javax.servlet.http.HttpServletRequest.getSession(false); %>` to get the already existing session that was created by the original session creating JSP file. This action helps prevent breaking session affinity on the initial loading of the frame pages.
- Update session data using only one frame. When using framesets, requests come into the HTTP server concurrently. Modifying session data within only one frame so that session changes are not overwritten by session changes in concurrent frameset is recommended.
- Avoid using multi-framed JSP files where the frames point to different Web applications. This action results in losing the session created by another Web application because the JSESSIONID cookie from the first Web application gets overwritten by the JSESSIONID created by the second Web application.
- **Secure all of the pages (not just some) when applying security to servlets or JSP files that use sessions with security integration enabled, .**

When it comes to security and sessions, it is all or nothing. It does not make sense to protect access to session state only part of the time. When security integration is enabled in the Session Management facility, all resources from which a session is created or accessed must be either secured or unsecured. You cannot mix secured and unsecured resources.

The problem with securing only a couple of pages is that sessions created in secured pages are created under the identity of the authenticated user. Only the same user can access sessions in other secured pages. To protect these sessions from use by unauthorized users, you cannot access these sessions from an unsecure page. When a request from an unsecure page occurs, access is denied and an `UnauthorizedSessionRequestException` error is thrown.

(`UnauthorizedSessionRequestException` is a runtime exception; it is logged for you.)

- **Use manual update and either the `sync()` method or time-based write in applications that read session data, and update infrequently.**

With `END_OF_SERVICE` as write frequency, when an application uses sessions and anytime data is read from or written to that session, the `LastAccess` time field updates. If database sessions are used, a new write to the database is produced. This activity is a performance hit that you can avoid using the Manual Update option and having the record written back to the database only when data values update, not on every read or write of the record.

To use manual update, turn it on in the Session Management Service. (See the tables above for location information.) Additionally, the application code must use the `com.ibm.websphere.servlet.session.IBMSession` class instead of the

generic HttpSession. Within the IBMSession object there is a method called sync(). This method tells the WebSphere Application Server to write the data in the session object to the database. This activity helps the developer to improve overall performance by having the session information persist only when necessary.

**Note:** An alternative to using the manual updates is to utilize the timed updates to persist data at different time intervals. This action provides similar results as the manual update scheme.

- Implement the following suggestions to achieve high performance:
  - If your applications do not change the session data frequently, use Manual Update and the sync() function (or timed interval update) to efficiently persist session information.
  - Keep the amount of data stored in the session as small as possible. With the ease of using sessions to hold data, sometimes too much data is stored in the session objects. Determine a proper balance of data storage and performance to effectively use sessions.
  - If using database sessions, use a dedicated database for the session database. Avoid using the application database. This helps to avoid contention for JDBC connections and allows for better database performance.
  - If using memory to memory sessions, define replicators only on the servers and have the client attach to server replicator.
  - If using memory to memory sessions, employ partitioning (either group or single replica) as your clusters grow in size and scaling decreases.
  - Verify that you have the latest e-fixes for the WebSphere Application Server.
- Utilize the following tools to help monitor session performance.
  - Run the com.ibm.servlet.personalization.sessiontracking.IBMTrackerDebug servlet. - To run this servlet, you must have the servlet invoker running in the Web application you want to run this from. Or, you can explicitly configure this servlet in the application you want to run.
  - Use the WebSphere Application Server Resource Analyzer which comes with WebSphere Application Server to monitor active sessions and statistics for the WebSphere Application Server environment.
  - Use database tracking tools such as "Monitoring" in DB2. (See the respective documentation for the database system used.)

---

## Managing HTTP sessions: Resources for learning:

Use the following links to find relevant supplemental information about HTTP sessions. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:


- [Best practices](#)
- [HTTP Session Persistence Best Practices](#)
- [Improving session persistence performance with DB2](#)

-  Persistent client state HTTP cookies specification

#### Programming instructions and examples

-  Java Servlet documentation, tutorials, and examples site

#### Programming specifications

-  Java Servlet 2.3 API specification download site
-  J2EE 1.3 specification download site



---

## Chapter 4. Using enterprise beans in applications

### Steps for this task

1. Design a J2EE application and the enterprise beans that it needs.  
See "Resources for learning" for links to design information that is specific to enterprise beans.
2. Develop any enterprise beans that your application will use.
3. Prepare for assembly. For your EJB 2.x-compliant entity beans, decide on an appropriate access intent policy.
4. Assemble the beans into one or more EJB modules. This includes "Securing enterprise bean applications" (not in this document).
5. (Assemble the modules) into a J2EE application.
6. For a given application server, update the EJB container configuration if needed for the application to be deployed, and determine if you want to "Setting the run time for batched commands" or "Setting the run time for deferred create" for container managed persistence.
7. Deploy the application in an application server.
8. Test the modules.
  - As needed, debug problems with the container.
  - Debug access and deployment problems.
9. (Assemble the production application.)
10. Deploy the application to a production environment.
11. Manage the application:
  - a. Manage installed EJB modules.  
After an application has been installed, you can manage its EJB modules individually through administrative console settings.
  - b. (Manage other aspects of the J2EE application.)
12. (Update the module and redeploy it.)
13. Tune the performance of the application. See "Best practices for developing enterprise beans".

---

## Enterprise beans

An enterprise bean is a Java component that can be combined with other resources to create J2EE applications. There are three types of enterprise beans, *entity* beans, *session* beans, and *message-driven* beans.

All beans reside in EJB containers, which provide an interface between the beans and the application server on which they reside.

Entity beans store permanent data. Entity beans with container-managed persistence (CMP) require connections to a form of persistent storage. This storage might be a database, an existing legacy application, a file, or other types of persistent storage. Entity beans with bean-managed persistence manage permanent data in whichever manner is defined in the bean code. This can include writing to databases or XML files, for example.

Session beans do not require database access, although they can obtain it indirectly as needed through entity beans. Session beans can also obtain direct access to databases (and other resources) through the use of resource references. Session beans can be either *stateful* or *stateless*.

New in the Enterprise JavaBeans (EJB) specification, version 2.0, message-driven beans enable asynchronous message servicing. The EJB container and a Java Message Service (JMS) provider work together to process messages. When a message arrives from another application component through JMS, the EJB container forwards it through an `onMessage()` call to a message-driven bean instance, which then processes the message. In other respects, message-driven beans are similar to stateless session beans.

Beans that require data access use *data sources*, which are administrative resources that define pools of connections to persistent storage mechanisms.

For more information about enterprise beans, see "Resources for learning."

---

## Developing enterprise beans

### Before you begin

Design a J2EE application and the enterprise beans that it needs.

- For general design information, see "Resources for learning."
- Before developing entity beans with container-managed persistence (CMP), read "Concurrency control."

There are two basic approaches to selecting tools for developing enterprise beans:

- You can use one of the available integrated development environments (IDEs). IDE tools automatically generate significant parts of the enterprise bean code and contain integrated tools for packaging and testing enterprise beans. The IBM WebSphere Application Developer product is the recommended IDE. For more information, see the documentation for that product.
- If you have decided to develop enterprise beans without an IDE, you need at least an ASCII text editor. You can also use a Java development tool that does not support enterprise bean development. You can then use tools available in the Java Software Development Kit (SDK) and in this product to assemble, test, and deploy the beans.

The following steps primarily support the second approach, development without an IDE.

### Steps for this task

1. If necessary, migrate any pre-existing code to the required version of the Enterprise JavaBeans (EJB) specification.
2. Write and compile the components of the enterprise bean.
  - At a minimum, an EJB 1.1 session bean requires a bean class, a home interface, and a remote interface. An EJB 1.1 entity bean requires a bean class, a primary-key class, a home interface, and a remote interface.
  - At a minimum, an EJB 2.0 session bean requires a bean class, a home or local home interface, and a remote or local interface. An EJB 2.0 entity bean requires a bean class, a primary-key class, a remote home or local home



interface, and a remote or local interface. The types of interfaces go together: If you implement a local interface, you must define a local home interface as well.

- Available only through EJB 2.0, a message-driven bean requires only a bean class.
3. For each entity bean, complete work to handle persistence operations.
- Create a database schema for the entity bean's persistent data.
    - For entity beans with container-managed persistence(CMP), you must store the bean's persistent data in one of the supported databases. The Application Assembly Tool automatically generates SQL code for creating database tables for CMP entity beans. If your CMP beans require complex database mappings, it is recommended that you use the IBM WebSphere Studio Application Developer product to generate code for the database tables.
    - For entity beans with bean-managed persistence (BMP), you can create the database and database table by using the database tools or use an existing database and database table.

For more information on creating databases and database tables, consult your database documentation.

- **(CMP entity beans for EJB 2.0 only)** (Define finder queries) with EJB Query Language (EJB QL).

With EJB QL, you define finders in terms of CMP fields and container-managed relationships, as follows:

- *Public* finders are visible in the bean's home interface. Implemented in the bean class, they return only remote interfaces and collection types.
  - *Private* finders, expressed as SELECT statements, are used only within the bean class. They can return both local and remote interfaces, dependent values, other CMP field types, and collection types.
- **(CMP entity beans for EJB 1.1 only: an IBM extension)** Create a finder helper interface for each CMP entity bean that contains specialized finder methods (other than the `findByPrimaryKey` method).

The following logic is required for each finder method (other than the `findByPrimaryKey` method) contained in the home interface of an entity bean with CMP:

- The logic must be defined in a public interface named `<i>Name</i>BeanFinderHelper`, where *Name* is the name of the enterprise bean (for example, `AccountBeanFinderHelper`).
- The logic must be contained in a String constant named `<i>findMethodName</i> WhereClause`, where *findMethodName* is the name of the finder method. The String constant can contain zero or more question marks (?) that are replaced from left to right with the value of the finder method's arguments when that method is called.

### What to do next

Assemble the beans in one or more EJB modules.

## Migrating enterprise bean code to the supported specification

Support for Version 2.0 of the Enterprise JavaBeans (EJB) specification is new for Version 5 of this product. Migration of enterprise beans deployed in Version 4.0.x

of this product is not generally necessary; Version 1.1 of the EJB specification is still supported. Follow these steps as appropriate for your application deployment.

#### Steps for this task

1. Modify enterprise bean code for changes in the specification.
  - For Version 1.0 beans, migrate at least to Version 1.1.
  - As stated previously, migration from Version 1.1 to Version 2.0 of the EJB specification is not required for redeployment on this version of the product. However, if your application requires the capabilities of Version 2.0, migrate your Version 1.1-compliant code.

**Note:** The EJB Version 2.0 specification mandates that prior to the EJB container's executing a *findByMethod* query, the state of all enterprise beans enlisted in the current transaction be synchronized with the persistent store. (This is so the query is performed against current data.) If Version 1.1 beans are reassembled into an EJB 2.0-compliant module, the EJB container synchronizes the state of Version 1.1 beans as well as that of Version 2.0 beans. As a result, you might notice some change in application behavior even though the application code for the Version 1.1 beans has not been changed.

2. Modify enterprise bean code for changes in deployment requirements.

If the enterprise beans were previously deployed in Version 3.0.x of this product, modify import statements to match standard package names. In Version 3.0.2.x, the following standard packages were present under nonstandard names:

```
javax.sql.*  
javax.transaction.*
```

Any code using WebSphere data sources, including BMP entity beans and session beans that access databases, must be modified.

3. You might have to modify code for some EJB 1.1-compliant modules that were not migrated to Version 2.0. Use the following information to help you decide.
  - Some stub classes for deployed enterprise beans have changed in the Java 2 SDK, Version 1.3.
  - The task of (generating deployment code) for enterprise beans changed significantly for EJB 1.1-compliant modules relative to EJB 1.0-compliant modules.
  - If the CMP beans write to databases with mixed-case table or column names and you used IBM VisualAge for Java, Version 3.5.x, to generate the original deployment code, you cannot simply reassemble the beans in this product. You must export the original EJB project from the VisualAge for Java product as an EJB 1.1 JAR. This preserves the metadata needed to generate the correct deployment code for mixed-case database tables and columns. For more information, see the documentation for the Deployment Tool for Enterprise JavaBeans(../..../deploytool/tasks/index.htm).

For detailed information about source and binary compatibility between deployed versions, see "Resources for learning."

4. (Reassemble) and (redeploy) all modules to incorporate migrated code.

### **Migrating enterprise bean code from Version 1.0 to Version 1.1**

The following information generally applies to any enterprise bean that currently complies with Version 1.0 of the Enterprise JavaBeans (EJB) specification. For more information about migrating code for beans produced with the IBM WebSphere

Studio Application Developer tool, see the documentation for that product. For more information about migrating code in general, see "Resources for learning."

#### Steps for this task

1. In session beans, replace all uses of `javax.jts.UserTransaction` with `javax.transaction.UserTransaction`. Entity beans may no longer use the `UserTransaction` interface at all.
2. In finder methods for entity beans, include `FinderException` in the throws clause.
3. Remove throws of `java.rmi.RemoteException`; throw `javax.ejb.EJBException` instead.  
However, continue to include `RemoteException` in the throws clause of home and remote interfaces as required by the use of Remote Method Invocation (RMI).
4. Remove uses of the `finalize()` method.
5. Replace calls to `getCallerIdentity()` with calls to `getCallerPrincipal()`.  
The use of `getCallerIdentity()` is deprecated.
6. Replace calls to `isCallerInRole(Identity)` with calls to `isCallerinRole (String)`.  
The use of `isCallerInRole(Identity)` and `java.security.Identity` is deprecated.
7. Replace calls to `getEnvironment()` in favor of JNDI lookup.  
As an example, change the following code:

```
String homeName =
    aLink.getEntityContext().getEnvironment().getProperty("TARGET_HOME_NAME");
if (homeName == null) homeName = "TARGET_HOME_NAME";
```

The updated code would look something like the following:

```
Context env = (Context)(new InitialContext()).lookup("java:comp/env");
String homeName = (String)env.lookup("ejb10-properties/TARGET_HOME_NAME");
```

8. In `ejbCreate` methods for an entity bean with container-managed persistence (CMP), return the bean's primary key class instead of `void`.
9. Add the `getHomeHandle()` method to home interfaces.  

```
public javax.ejb.HomeHandle getHomeHandle() {return null;}
```

#### What to do next

Consider enhancements to match the following changes in the specification:

- Primary keys for entity beans can be of type `java.lang.String`.
- Finder methods for entity beans return `java.util.Collection`.
- Check the format of any JNDI names being used. Local name spaces are also supported.
- Security is defined by role, and isolation levels are defined at the method level rather than at the bean level.

### **Migrating enterprise bean code from Version 1.1 to Version 2.0**

Enterprise JavaBeans (EJB) Version 2.0-compliant beans may be assembled only in an EJB 2.0-compliant module, although an EJB 2.0-compliant module can contain a mixture of Version 1.x and Version 2.0 beans.

The EJB Version 2.0 specification mandates that prior to the EJB container's executing a *findByMethod* query, the state of all enterprise beans enlisted in the current transaction be synchronized with the persistent store. (This is so the query is performed against current data.) If Version 1.1 beans are reassembled into an EJB

2.0-compliant module, the EJB container synchronizes the state of Version 1.1 beans as well as that of Version 2.0 beans. As a result, you might notice some change in application behavior even though the application code for the Version 1.1 beans has not been changed.

The following information generally applies to any enterprise bean that currently complies with Version 1.1 of the EJB specification. For more information about migrating code for beans produced with the IBM WebSphere Studio Application Developer tool, see the documentation for that product. For more information about migrating code in general, see "Resources for learning."

#### Steps for this task

1. In beans with container-managed persistence (CMP) version 1.x, replace each CMP field with abstract get and set methods.  
In doing so, you must make each bean class abstract.
2. In beans with CMP version 1.x, change all occurrences of this.`<i>field</i> = <i>value</i>` to `set<i>Field</i>(<i>value</i>)`.
3. In each CMP bean, create abstract get and set methods for the primary key.
4. In beans with CMP version 1.x, create an EJB Query Language statement for each finder method.
5. In finder methods for beans with CMP version 1.x, return `java.util.Collection` instead of `java.util.Enumeration`.
6. Update handling of non-application exceptions.
  - To report non-application exceptions, throw `javax.ejb.EJBException` instead of `java.rmi.RemoteException`.
  - Modify rollback behavior as needed: In EJB versions 1.1 and 2.0, all non-application exceptions thrown by the bean instance result in the rollback of the transaction in which the instance is running; the instance is discarded. In EJB 1.0, the container does not roll back the transaction or discard the instance if it throws `java.rmi.RemoteException`.
7. Update rollback behavior as the result of application exceptions.
  - In EJB versions 1.1 and 2.0, an application exception does not cause the EJB container to automatically roll back a transaction.
  - In EJB Version 1.1, the container performs the rollback only if the instance has called `setRollbackOnly()` on its `EJBContext` object.
  - In EJB Version 1.0, the container is required to roll back a transaction when an application exception is passed through a transaction boundary started by the container.

## **WebSphere extensions to the Enterprise JavaBeans specification**

This article outlines extensions to the Enterprise JavaBeans (EJB) specification that IBM provides with this product:

### **Inheritance in enterprise beans**

In the Java language, *inheritance* is the creation of a new class from an existing class or a new interface from an existing interface. This product supports two forms of inheritance: standard class inheritance and EJB inheritance.

In standard class inheritance, the home interface, remote interface, or enterprise bean class inherits properties and methods from base classes that are not themselves enterprise bean classes or interfaces.

By contrast in enterprise bean inheritance, an enterprise bean inherits properties (such as container-managed persistence (CMP) fields and container-managed relationship (CMR) fields), methods, and method-level control descriptor attributes from another enterprise bean.

For more information, see the documentation for the IBM WebSphere Studio Application Developer product.

### **Optimistic concurrency control for container-managed persistence**

This product supports optimistic concurrency control of data access.

### **Access intents for EJB persistence**

#### **5.0.2**

This product supports the application of named data-access policies.

### **Performance enhancements**

Through the lifetime-in-cache settings, this product provides a way for you to improve performance for beans that are only occasionally updated. For more information, see "Entity bean assembly settings."

Some enterprise beans created with the IBM WebSphere Studio Application Developer product can utilize *read-ahead* for loading a bean and its related beans in a single database operation. An entire object graph or any part of the graph can be preloaded by configuring a finder method to use read-ahead.

### **Assembly and deployment extensions**

This product supports IBM extensions of assembly and deployment options. IBM extensions are clearly marked in reference topics for assembly settings.

## **Best practices for developing enterprise beans**

Use the following guidelines when designing and developing enterprise beans:

- Use a stateless session bean to act as the entry point for business logic. For more information about using session facades, see "Resources for learning."
- Entity beans should use container-managed persistence.
- In an Enterprise JavaBeans (EJB) Version 2.0 environment, use local interfaces to improve communication between enterprise beans in the same Java virtual machine.

Local calls avoid the overhead of RMI/IIOP and use pass-by-reference semantics instead of pass-by-value. For each call, the caller and callee beans share the state of arguments. EJB 2.0 beans can have both a local and remote interface but more typically have one or the other.

- For communicating with remote clients, provide remote and remote home interfaces. For communicating with local clients like servlets, entity beans, and message-driven beans, provide local and local home interfaces.

## Batch commands for container managed persistence

From JDBC 2.0 on, *PreparedStatement* objects can maintain a list of commands that can be submitted together as a batch. Instead of multiple database round trips, there can be only one database round trip for all the batched persistence requests.

The WebSphere Application Server version 5.0.2 enables you to take advantage of this. You can turn this option on from the EJB CMP side. When you choose this option, the run time defers *ejbStore/ejbCreate/ejbRemove* or the equivalent database persistence requests (insert/update/delete) until they are needed. This can be at the end of the transaction, or when a flush is needed for finders related to this EJB type. When the persistence operation finally happens, run time accumulates the database requests and uses JDBC *PreparedStatement* batch operation to make a single JDBC call for multiple rows of the same operation.

## Setting the run time for batched commands

### Steps for this task

1. Open the administrative console.
2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.
5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with *Dcom.ibm.ws.pm.batch=true*.

## Deferred Create for container managed persistence

The specification for Enterprise Java Beans (EJB) 2.x states that for Container Managed Persistence (CMP) during the *ejbCreate*, the container can create the representation of the entity in the database immediately, or defer it to a later time.

The WebSphere Application Server version 5.0.2 enables you to take advantage of this specification. You can turn this option on from the EJB CMP side. When you choose this option, the runtime defers *ejbCreate* (or the equivalent database persistence request) until it is needed. This can be at the end of the transaction, or when a flush is needed for finders related to this EJB type. By doing this you can reduce two round trips for the newly created entity (insert and update) to one (insert).

## Setting the run time for deferred create

### Steps for this task

1. Open the administrative console.
2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.
5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with *Dcom.ibm.ws.pm.deferredcreate=true*.

## Explicit Invalidation in the Persistence Manager Cache

Container managed persistence (CMP) entity beans can be configured as *long-lifetime* beans. A long-lifetime bean is one that is configured with *Lifetime In Cache Usage* equal to a value other than the default **OFF** (refer to "Entity bean assembly settings"). A value other than **OFF** means that data for this bean is cached beyond the end of the transaction in which the bean was obtained by a

finder or other method. The *Lifetime In Cache Usage* and *Lifetime In Cache* values control the basic length of time the cached data remains valid. When the specified time runs out, the cached data is no longer valid. See the *LifetimeInCache* help sections of the Application Assembly Tool (AAT) for more details.

However, there is also an API that lets the client application code explicitly invalidate the cached data of a bean on demand, superceding the basic lifetime of the cache data as controlled by the *Lifetime In Cache Usage* and *Lifetime In Cache* settings. This is useful where an application that does not use CMP beans modifies the data that underlies a CMP bean (for example, it updates a database table to which a CMP bean is mapped). Such an application can inform WebSphere Application Server that any cached version of this bean data is **stale** and no longer matches what is in the database. The data should be invalidated (in essence, discarded). For CMP beans that cannot tolerate stale data, this is an important feature.

Because the PM Cache Invalidation mechanism does consume resources in exchange for its benefits, it is not enabled by default. To enable it refer to "Setting Persistence Manager Cache Invalidation".

### **Example: Explicit Invalidation in the Persistence Manager Cache Usage Scenario**

The scenario of use for this feature begins with configuring one or more bean types to be long-lifetime beans (see "Explicit Invalidation in the Persistence Manager Cache", and configuring the necessary Java Message Service (JMS) resources (described below). Once this is done, the server is started. The scenario continues as follows:

1. Assume that a CMP entity bean of type *Department* has been configured to be a long-lifetime bean.
2. Transaction 1 begins. Application code looks up *Department's* home and calls a finder method (such as *findByPrimaryKey("dept01")*). As this is the first finder to return *Department dept01*, a trip is made to the database to obtain the data. Transaction 1 ends.
3. Transaction 2 begins. Application code calls *findByPrimaryKey("dept01")* again. Because this is not the first finder to return *Department dept01*, we get a cache hit and no database trip is made. So far this is current WebSphere Application Server behavior for long-lifetime beans. Transaction 2 ends.
4. Another application, which does not use the *Department* CMP bean, is executed. This application might or might not be run on the WebSphere Application Server; it could be a legacy application. The application updates the database table that is mapped to the *Department* bean, altering the row for *dept01*. For example, the *budget* column in the table (mapped to a Java double CMP attribute in the *Department* bean) is changed from \$10,000.00 to \$50,000.00. This application was designed to cooperate with WebSphere Application Server. After performing the update, the application sends an invalidate request message to invalidate the *Department* bean *dept01*.
5. Transaction 3 begins. Application code looks up *Department's* home and calls a finder method (such as *findByPrimaryKey("dept01")*). Because this is the first finder after *Department dept01* is invalidated, a new database trip is made to obtain the data. Transaction 3 ends.

### **Persistence Manager cache invalidation API**

The PM cache invalidation API is in the form of a JMS message that the client sends to a specially-named JMS topic using a connection from a specifically named JMS *TopicConnectionFactory*. The JMS message must be an *ObjectMessage* created by the client. The client code creates a *PMCacheInvalidationRequest* object that describes the bean data to invalidate. Client code places the *PMCacheInvalidationRequest* object in the *ObjectMessage* and publishes the *ObjectMessage* (for further details on the JMS objects and terms used here, please see the Java Message Service documentation).

The public class *PMCacheInvalidationRequest* is central to the API, so we include a portion of its code here for illustration purposes (if you see any differences between this illustration and the actual class, the class is to be considered correct):

```
package com.ibm.websphere.ejbpersistence;

/**
 *An instance of this class represents a request to invalidate one or more
 *CMP beans in the PMcache. When an invalidate occurs, cached data for this
 *bean is removed from the cache; the next time an application tries to find
 *this bean, a fresh copy of the bean data is obtained from the data store.
 *
 *The ability to invalidate a bean means that a CMP bean may be configured
 *as a long-lifetime bean and thus be cached across transactions for much
 *greater performance on future attempts to find this bean. Yet when some
 *outside mechanism updates the bean data, sending an invalidation request
 *will remove stale data from the PMcache so applications do not behave falsely
 *based on stale data.
 */
public class PMCacheInvalidationRequest implements Serializable {
    . . .

    /**
     * Constructor used to invalidate a single bean
     * @param beanHomeJNDIName the JNDI name of the bean home. This is the same value
     * used to look up the bean home prior to calling findByPrimaryKey, for example.
     * @param beanKey the primary key of the bean to be invalidated. The actual
     * object type must be the primary key type for this bean type.
     */
    public PMCacheInvalidationRequest(String beanHomeJNDIName, Object beanKey)
        throws IOException {
        . . .
    }

    /**
     * Constructor used to invalidate a Collection of beans
     * @param beanHomeJNDIName java.lang.String the JNDI name of the bean home.
     * This is the same value used to look up the bean home prior to calling
     * findByPrimaryKey, for example.
     * @param beanKeys a Collection of the primary keys of the beans to be
     * invalidated. The actual type of each object in the Collection must be the
     * primary key type for this bean type.
     */
    public PMCacheInvalidationRequest(String beanHomeJNDIName, Collection beanKeys)
        throws IOException {
        . . .
    }

    /**
     * Constructor used to invalidate all beans of a given type
     * @param beanHomeJNDIName java.lang.String the JNDI name of the bean home.
     * This is the same value used to look up the bean home prior to calling
     * findByPrimaryKey, for example.
     */
    public PMCacheInvalidationRequest(String beanHomeJNDIName) {
```



```

    . . .
}
}

```

If the client wants to perform the invalidation in a synchronous way, it can opt to receive an acknowledgement JMS message when the invalidation is complete. To ask for an acknowledgement (ACK) message, the client sets a *Topic* of its own choosing in the *JMSReplyTo* field of the *ObjectMessage* for the invalidation request (see JMS documentation for further details). The client then waits (using the *receive()* method of JMS) on receipt of the acknowledgement message before continuing execution.

An ACK message enables the caller to insure there is not even a brief (seconds or less) window during which PM cache data is stale. The sending of an acknowledgement for each request does, of course, take a bit more time and so is recommended to be used only when needed.

The JMS resources used to make an invalidation request (*TopicConnectionFactory*, *TopicDestination*, and so forth) must be configured by the user (using the Administration console or other method) if they want to use PM Cache Invalidation. In this way the user can choose whichever JMS provider they prefer (as long as it supports pub-sub). The names that must be used for these resources are defined as part of the API, and use names unique to the WebSphere Application Server namespace to avoid name conflict with customer JMS resources.

The following are the names that must be used when the user configures the JMS resources (shown as Java constants for clarity):

```

// The JNDI name of the TopicConnectionFactory
private static final String topicConnectionFactoryJNDIName =
    "com.ibm.websphere.ejbpersistence.InvalidatetCF";
// The JNDI name of the TopicDestination
private static final String topicDestinationJNDIName =
    "com.ibm.websphere.ejbpersistence.invalidate";
// The Topic name (part of the TopicDestination)
private static final String topicString =
    "com.ibm.websphere.ejbpersistence.invalidate";

```

Here are examples of how these constants can be used in client code:

```

// Look up the TopicConnectionFactory...
InitialContext ic = new InitialContext();
TopicConnectionFactory topicConnectionFactory =
    (TopicConnectionFactory) ic.lookup(topicConnectionFactoryJNDIName);
...
// Look up the Topic
Topic topic = (Topic) ic.lookup(topicDestinationJNDIName);

```

Note that JMS messages can be sent not only from J2EE application code (for example, a *SessionBean* or *BMP* entity bean method) but also from non-J2EE applications if your chosen JMS provider allows for this. For example, the IBM MQ provider, available in WebSphere Application Server as the **Embedded Messaging** feature (selectable during installation), supports the use of MQ classes (or structures in other languages) to create a topic connection and topic that are compatible with the *TopicConnectionFactory* and *TopicDestination* you configure using WebSphere Application Server Application Console.

## Setting Persistence Manager Cache Invalidation

### Steps for this task

1. Open the administrative console.

2. Select **Servers**.
3. Select **Application Servers**.
4. Select the server you want to configure.
5. In the Additional Properties area, select **Process Definition**.
6. In the Additional Properties area, select **Java Virtual Machine**.
7. Update the **Generic JVM arguments** with  
*-Dcom.ibm.ws.ejbpersistence.cacheinvalidation=true.*

---

## Using access intent policies

### 5.0.2

You can use access intent policies to help the product run-time environment manage various aspects of Enterprise JavaBeans (EJB) persistence. You apply access intent policies to EJB Version 2.0 entity beans and their methods by using the Application Assembly Tool. This product provides a set of default access intent policies. You can also create your own custom policies.

#### Steps for this task

1. Apply default access intent to CMP entity beans.  
For more information, see Entity bean assembly settings.
2. Apply access intent policies to methods of CMP entity beans.
3. **(Optional)** (Create a custom access intent policy) by using the Application Assembly Tool.
4. **(Optional)** Apply access intent policies to BMP entity bean methods by using the AccessIntent API.
5. **(Optional)** Apply multiple access intent policies to methods by using (application profiling).

## Access intent policies

An access intent policy is a named set of properties (access intents) that governs data access for Enterprise JavaBeans (EJB) persistence. You can assign policies to an entity bean and to individual methods on an entity bean's home, remote, or local interfaces during assembly. If you have the WebSphere Application Server Enterprise product installed, you can assign these during development as well. Access intents are settable only within EJB Version 2.x-compliant modules for entity beans with bean-managed persistence or with CMP Version 2.x.

This product supplies a number of access intent policies that specify permutations of read intent and concurrency control; the pessimistic/update policy can be qualified further. The selected policy determines the appropriate isolation level and locking strategy used by the run-time environment.

Access intent policies are specifically designed to supplant the use of isolation level and access intent method-level modifiers found in the extended deployment descriptor for EJB version 1.1 enterprise beans. You cannot specify isolation level and read-only modifiers for EJB version 2.0 enterprise beans.

### 5.0.2

Access intent can be controlled in a more precise way by using either application profiling or by using method-level access intent policies. Application profiling is

only available in the WebSphere Application Server Enterprise product. Method-level access intent policies are named and defined at the module level. A module can have one or many such policies. Policies are assigned, and apply, to individual methods of the declared interfaces of entity beans and their associated home interfaces. A policy is acted upon by either the combination of the EJB container and persistence manager (for entity beans with container-managed persistence) or directly by entity beans with bean-managed persistence.

For entity beans that are backed by tables with nullable columns, use an optimistic policy with caution. Nullable columns are automatically excluded from overqualified updates at deployment time; concurrent changes to a nullable field might result in lost updates. When used with the IBM WebSphere Studio Application Developer product, this product provides support for selecting a subset of the nonnullable columns that are to be reflected in the overqualified update statement that is generated in the deployment code to support optimistic policies.

### 5.0.2

An entity that is configured with a read-only policy that causes a bean to be activated can cause problems if updates are attempted within the same transaction. Those changes will not be committed, and an exception will be thrown because data integrity might be compromised.

### Concurrency control

Concurrency control is the management of contention for data resources. A concurrency control scheme is considered *pessimistic* when it locks a given resource early in the data-access transaction and does not release it until the transaction is closed. A concurrency control scheme is considered *optimistic* when locks are acquired and released over a very short period of time at the end of a transaction.

The objective of optimistic concurrency is to minimize the time over which a given resource would be unavailable for use by other transactions. This is especially important with long-running transactions, which under a pessimistic scheme would lock up a resource for unacceptably long periods of time.

Under an optimistic scheme, locks are obtained immediately before a read operation and released immediately afterwards. Update locks are obtained immediately before an update operation and held until the end of the transaction.

To enable optimistic concurrency, this product uses an *overqualified update scheme* to test whether the underlying data source has been updated by another transaction since the beginning of the current transaction. With this scheme, the columns marked for update and their original values are added explicitly through a WHERE clause in the UPDATE statement so that the statement fails if the underlying column values have been changed. As a result, this scheme can provide column-level concurrency control; pessimistic schemes can control concurrency at the row level only.

Optimistic schemes typically perform this type of test only at the end of a transaction. If the underlying columns have not been updated since the beginning of the transaction, pending updates to container-managed persistence fields are committed and the locks are released. If locks cannot be acquired or if some other transaction has updated the columns since the beginning of the current transaction, the transaction is rolled back: All work performed within the transaction is lost.

Pessimistic and optimistic concurrency schemes require different transaction isolation levels. Enterprise beans that participate in the same transaction and require different concurrency control schemes cannot operate on the same underlying data connection.

Whether or not to use optimistic concurrency depends on the type of transaction. Transactions with a high penalty for failure might be better managed with a pessimistic scheme. (A high-penalty transaction is one for which recovery would be risky or resource-intensive.) For low-penalty transactions, it is often worth the risk of failure to gain efficiency through the use of an optimistic scheme. In general, optimistic concurrency is more efficient when update collisions are expected to be infrequent; pessimistic concurrency is more efficient when update collisions are expected to occur often.

### Read-ahead hints

Read-ahead schemes enable applications to minimize the number of database roundtrips by retrieving a working set of container-managed persistence (CMP) beans for the transaction within one query. Read-ahead involves activating the requested CMP beans and caching the data for their related beans, which ensures that data is present for the beans that are most likely to be needed next by an application. A *read-ahead hint* is a canonical representation of the related beans that are to be read. It is associated with the *findByPrimaryKey* method for the requested bean type, which must be an EJB 2.x-compliant CMP entity bean.

#### 5.0.2

Read-ahead hints can be set only using the WebSphere Application Server Enterprise assembly tool or through the Add Access Intent wizard of the IBM WebSphere Studio Application Developer product.

#### 5.0.2

Read-ahead is only supported for access intent policies that can be applied by the backend against which the application is deployed. Otherwise, the read-ahead hint is disregarded.

#### 5.0.2

Currently, only *findByPrimaryKey* methods can have read-ahead hints. Only beans related to the requested beans by a container-managed relationship (CMR), either directly or indirectly through other beans, can be read ahead. Beans that use EJB inheritance should not be used in a read-ahead hint.

A read-ahead hint takes the form of a character string. You do not have to provide the string; the wizard generates it for you based on CMRs defined for the bean. The following example is provided as supplemental information only.

Suppose a CMP bean type A has a finder method that returns instances of bean A. A read-ahead hint for this method is specified using the following notation:  
`<i>Re1B</i>.<i>Re1C</i>; <i>Re1D</i>`

Interpret the preceding notation as follows:

- Bean type A has a CMR with bean types B and D.
- Bean type B has a CMR with bean type C.

For each bean of type A that is retrieved from the database, its directly-related B and D beans and its indirectly-related C beans are also retrieved. The order of the retrieved bean data columns in each row of the result set is the same as their order in the read-ahead hint: an A bean, a B bean (or null), a C bean (or null), a D bean (or null). For hints in which the same relationship is mentioned more than once (for example, `<i>RelB</i>.<i>RelC</i>;<i>RelB</i>.<i>RelC</i>`), a bean's data columns appear only once, at the position it first appears in the hint.

The tokens shown in the notation (*RelB* and so on) must be CMR field names for the relationships as defined in the deployment descriptor for the bean. In indirect relationships such as `<i>RelB</i>.<i>RelC</i>`, *RelC* is a CMR field name defined in the deployment descriptor for bean type B.

A single read-ahead hint cannot refer to the same bean type in more than one relationship. For example, if a Department bean has a relationship *employees* with the Employee bean and also has a relationship *manager* with the Employee bean, the read-ahead hint cannot specify both *employees* and *manager*.

For more information about how to set read-ahead hints, see the documentation for the Websphere Studio Application Developer product.

## Access intent service

Access intent is a WebSphere Application Server run-time service that enables you to more precisely manage an application's persistence. The access intent service defines a set of declarative annotations used by the Enterprise JavaBeans (EJB) container and its agents to make performance optimizations for entity bean access. These annotations are organized into sets called *access intent policies*.

Access intent policies contain a set of annotations considered as hints by the EJB container and its agents. Most access intent policies are hints representing high-level abstractions that can be mapped to a specific backend resource manager. It is the responsibility of the EJB persistence machinery to ensure the necessary concurrency control, connection, and cache management when carrying out the persistence details. The EJB persistence manager can use access intent hints to make better performance decisions when carrying out its assigned task. A smaller number of access intents are hints to the EJB container, influencing the management of EJB collections.

You can apply access intent policies to methods within the scope of an EJB module, in which case the policy becomes the default access intent for all requests upon the configured methods.

You can also apply access intent policies to methods within the scope of application profiles. Consequently, you can configure methods with multiple and opposing access intent policies. The application profiling documentation explains in more detail how to configure an application to apply a particular access intent policy to a method for one request, then apply another access intent policy to the same method for a different request.

### Access intent with BMP entity beans

Access intent's declarative functionality provides great power to you as a CMP entity bean developer. You can provide hints on how WebSphere Application Server is to manage the details of persistence without having to explicitly manage any of the persistence logic from within the application.

There are situations, however, in which you might need to develop BMP entity beans. Because the only meaningful difference between BMP and CMP components is who provides the persistence logic, BMP entity beans should be able to leverage access intent hints just as WebSphere Application Server does on behalf of CMP entity beans. BMP entity beans that use the access intent service participate in application profiling; that is, the value of the access intent attributes can differ from request to request, allowing the BMP entity bean to seamlessly modify its persistence strategy.

You can apply access intent policies to BMP entity bean methods as well as CMP entity bean methods. Because access intent hints are not contractual in nature, there is no obligation for a BMP entity bean to exploit them. BMP entity beans are expected to use only those access intent attributes that are important to that particular bean.

The current access intent policy is bound into the `java:comp` namespace for a particular BMP entity bean. That policy is current only for the duration of the method call during which the access intent policy was retrieved. In a typical scenario, you would cache the access type during invocation of the `ejbLoad()` method so that appropriate actions can be taken during invocation of the `ejbStore()` method.

## Access intent design considerations

Use the access intent service to solve clear performance problems. Identify usage patterns that lead to poor application performance and apply appropriate access intent policies.

Refrain from over-tuning an application. You can introduce errors by incorrectly using the access intent service. For example, misuse of the `wsPessimisticUpdate-NoCollision` policy can result in lost updates; inappropriately setting the collection increment value can introduce performance issues; and problem determination is more difficult when an application is confusingly configured with multiple access intent policies. Clarity and simplicity should be your guiding principles when using the access intent service. This is even more important when applying access intent policies within the scope of application profiles (a feature of WebSphere Application Server Enterprise).

Even though access intent policies can be configured on any method of an entity bean, some attributes of a policy can only be leveraged by the run-time environment under certain conditions. For example, concurrency and access intent are only used for CMP entity beans when the `ejbLoad()` method is driven to open a connection to, and read data from, a given resource; that data is cached and used to drive the proper queries during invocation of the `ejbStore()` method. Read-ahead hints are only used during the execution of a finder for a bean. Finally, the collection increment and resource manager prefetch increment are only used on multi-object finders. Configuring policies on methods that will not use the policy is not an error (only certain attributes of any policy are used, even when the policy is appropriately applied to a method). However, configuring policies unnecessarily throughout an application obscures the design of the application and complicates the maintenance of the application.

## Applying access intent policies to methods

You apply an access intent policy to a method, or set of methods, in an application's entity beans through the Application Assembly Tool (AAT).

### Steps for this task

1. (Start the AAT.)
2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, select **File > Open**, then select the EAR file.
3. Select **EJB Modules > *moduleName* > Access Intent**.
4. To configure a new access intent policy, right-click and select **New**.
5. On the **New Access Intent** panel, specify a name and a description.  
These attributes are provided as a convenience to the developer and are not used at run time.
6. To select the methods to which the access intent policy should apply, click **Add** beside the Methods table.
7. From the **Applied access intent** list, select an access intent policy.
8. **(Optional)** To override an attribute defined in the applied policy, click **Add** beside the Access intent attribute overrides table.
9. Click **OK** to exit the New Access Intent panel.
10. Save your configuration by selecting **File > Save**.

## Using the AccessIntent API

This task describes how to programmatically retrieve and call the AccessIntent API during the execution of BMP entity bean methods.

### Steps for this task

1. Look up the current access intent in the namespace.

For example:

```
InitialContext ic = new InitialContext();
AccessIntent ai = ic.lookup("java:comp/websphere/AppProfile/AccessIntent");
```

2. Call the necessary get() methods.

For example:

```
int concurrency = ai.getConcurrencyControl();
int accessType = ai.getAccessType();
if ( (concurrency == AccessIntent.CONCURRENCY_CONTROL_PESSIMISTIC)
    && (accessType == AccessIntent.ACCESS_TYPE_UPDATE) ) {
    boolean exclusive = ai.getPessimisticUpdateHintExclusive();
    // . . .
}
// . . .
```

### Results

**Note:** The access intent object reference retrieved from the java:comp lookup is current for the duration of the method in which the reference was looked up. Depending on how you configured the application profile, subsequent calls of the same method might not retrieve the same access intent reference. You can only look up the object reference during the call of a BMP entity bean's method; the reference does not exist during a request on a CMP entity bean. Therefore, access intent object references should not be cached beyond, or used outside of, the scope of the execution of any given BMP method.

### **AccessIntent interface**

The AccessIntent interface is available to BMP entity beans by the following JNDI lookup:

```
java:comp/websphere/AppProfile/AccessIntent
```

## AccessIntent interface

```
package com.ibm.websphere.appprofile.accessintent;

/**
 * This interface defines the essential access intents
 * available at runtime.
 */
public interface AccessIntent {

    /**
     * Returns the concurrency control intent, which indicates
     * the application prefers either pessimistic or optimistic
     * concurrency control when accessing the current component
     * in the context of the current transaction.
     */
    public int getConcurrencyControl();
    public final int CONCURRENCY_CONTROL_PESSIMISTIC = 1;
    public final int CONCURRENCY_CONTROL_OPTIMISTIC = 2;

    /**
     * Returns access type intent, which indicates the application
     * intends either update or read access of the current component
     * in the context of the current transaction.
     */
    public int getAccessType();
    public final int ACCESS_TYPE_UPDATE= 1;
    public final int ACCESS_TYPE_READ = 2;

    /**
     * Returns a boolean where true indicates that the runtime should
     * assume that there will be no collision on retrieved rows.
     */
    public boolean getPessimisticUpdateHintNoCollision();

    /**
     * Returns a boolean where true indicates that the runtime should
     * assume that there will be collisions on retrieved rows.
     */
    public boolean getPessimisticUpdateHintExclusive();

    /**
     * Returns the collection access intent, which indicates the
     * application intends to access the objects returned by the
     * currently executing finder in either serial or random fashion.
     */
    public int getCollectionAccess();
    public final int COLLECTION_ACCESS_RANDOM = 1;
    public final int COLLECTION_ACCESS_SERIAL = 2;

    /**
     * Returns the collection scope, which indicates the maximum
     * lifespan of a lazy collection.
     */
    public int getCollectionScope();
    public final int COLLECTION_SCOPE_TRANSACTION = 1;
    public final int COLLECTION_SCOPE_ACTIVITYSESSION = 2;
    public final int COLLECTION_SCOPE_TIMEOUT = 3;

    /**
     * Returns the timeout value in seconds when collectionScope is Timeout.
     */
    public int getCollectionTimeout();

    /**
     * Returns the number of elements the application requests be contained
     * in each segment of the element collection returned by the currently
     * executing finder.
     */
}
```



```

*/
public int getCollectionIncrement();

/**
 * Returns the ReadAheadHint requested by the application for the currently
 * executing finder.
 */
public ReadAheadHint getReadAheadHint();

/**
 * Returns the number of elements the application requests be contained in
 * each segment of a query made on a database.
 */
public int getResourceManagerPreFetchIncrement();
}

```

## Access intent exceptions

The following exceptions are thrown in response to the application of access intent policies:

### **com.ibm.ws.ejbpersistence.utilpm.PersistenceManagerException**

If the method that drives the `ejbLoad()` method is configured to be read-only but updates are then made within the transaction that loaded the bean's state, an exception is thrown during invocation of the `ejbStore()` method, and the transaction is rolled back. Likewise, the `ejbRemove()` method cannot succeed in a transaction that is set as read-only. If an update hint is applied to methods of entity beans with bean-managed persistence, the same behavior and exception results. The forwarded exception object contains the message string `PMGR1103E: update instance level read only bean <i>beanName</i>`

This exception is also thrown if the applied access intent policy cannot be honored because a finder, `ejbSelect`, or container-managed relationship (CMR) accessor method returns an inherently read-only result. The forwarded exception object contains the message string `PMGR1001: No such DataAccessSpec - <i>methodName</i>`

The most common occurrence of this error is when a custom finder that contains a read-only EJB Query Language (EJB QL) statement is called with an applied access intent of `wsPessimisticUpdate` or `wsPessimisticUpdate-Exclusive`. These policies require the use of a `FOR UPDATE` clause on the SQL `SELECT` statement to be executed, but a read-only query cannot support `FOR UPDATE`. Other examples of read-only queries include joins; the use of `ORDER BY`, `GROUP BY`, and `DISTINCT` keywords.

To eliminate the exception, edit the EJB query so that it does not return an inherently read-only result or change the access intent policy being applied.

- If an update access is required, change the applied access intent setting to `wsPessimisticUpdate-WeakestLockAtLoad` or `wsOptimisticUpdate`.
- If update access is not truly required, use `wsPessimisticRead` or `wsOptimisticRead`.
- If connection sharing between entity beans is required, use `wsPessimisticUpdate-WeakestLockAtLoad` or `wsPessimisticRead`.

### **com.ibm.websphere.ejb.container.CollectionCannotBeFurtherAccessed**

If a lazy collection is driven after it is no longer in scope, and beyond what has already been locally buffered, a `CollectionCannotBeFurtherAccessed` exception is thrown.

### **com.ibm.ws.exception.RuntimeWarning**

If an application is configured incorrectly, a run-time warning exception is thrown as the application starts; startup is ended. You can validate an application's configuration by choosing the verify function in the WebSphere Application Assembly Tool. Some examples of misconfiguration include:

- A method configured with two different access intent policies
- A method configured with an undefined access intent policy

### **javax.ejb.NoSuchEntityException**

#### **5.0.2**

If an update fails under optimistic concurrency because fields changed within another transaction between load and store requests, a `NoSuchEntityException` is raised and the commit fails.

## **Access intent assembly settings**

Access intent policies contain data-access settings for use by the persistence manager. Default access intent policies are configured on the entity bean. Optionally, you can associate access intent policies with one or more methods.

These settings are applicable only for EJB 2.x-compliant entity beans that are packaged in EJB 2.x-compliant modules. Connection sharing between beans with bean-managed persistence and those with container-managed persistence is possible if they all use the same access intent policy.

### **Name**

Specifies a name for a mapping between an access intent policy and one or more methods.

### **Description**

Contains text that describes the mapping.

### **Methods - Name**

Specifies the name of an enterprise bean method, or the asterisk character (\*). The asterisk is used to denote all of the methods of an enterprise bean's remote and home interfaces.

### **Methods - Enterprise bean**

Specifies which enterprise bean contains the methods indicated in the Name setting.

### **Methods - Type**

Used to distinguish between a method with the same signature that is defined in both the home and remote interface. Use `Unspecified` if an access intent policy applies to all methods of the bean.

Data type String

Range Valid values are Home, Remote, Local, LocalHome or Unspecified

### **Methods - Parameters**

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

## Applied access intent

Specifies how the container must manage data access for persistence. Configurable both as a default access intent for an entity and as part of a method-level access intent policy.

|           |  |
|-----------|--|
| Data type | String   |
| Default   | <code>wsPessimisticUpdate-WeakestLockAtLoad</code> .<br>With Oracle, this is the same as <code>wsPessimisticUpdate</code> .  |
| Range     | Valid settings are <code>wsPessimisticUpdate</code> , <code>wsPessimisticUpdate-NoCollision</code> , <code>wsPessimisticUpdate-Exclusive</code> , <code>wsPessimisticUpdate-WeakestLockAtLoad</code> , <code>wsPessimisticRead</code> , <code>wsOptimisticUpdate</code> , or <code>wsOptimisticRead</code> . Only <code>wsPessimisticRead</code> and <code>wsOptimisticRead</code> are valid when class-level caching is enabled in the EJB container. |

This product supports lazy collections. For each segment of a collection, iterating through the collection (`next()`) does not trigger a remote method call to retrieve the next remote reference. Two policies (`wsPessimisticUpdate` and `wsPessimisticUpdate-Exclusive`) are extremely lazy; the collection increment size is set to 1 to avoid overlocking the application. The other policies have a collection increment size of 25.

### 5.0.2

If an entity is not configured with an access intent policy, the run-time environment typically uses `wsPessimisticUpdate-WeakestLockAtLoad` by default. If, however, the `Lifetime in cache` property is set on the bean, the default value of **Applied access intent** is `wsOptimisticRead`; updates are not permitted.

Additional information about valid settings follows:

| Profile name   | Concurrency control | Access type | Transaction isolation                                     |
|--|---------------------|-------------|---|
| <code>wsPessimisticRead</code><br>(Note 1)                     | pessimistic         | read        | For Oracle, read committed.<br>Otherwise, repeatable read |
| <code>wsPessimisticUpdate</code><br>(Note 2)                   | pessimistic         | update      | For Oracle, read committed.<br>Otherwise, repeatable read |
| <code>wsPessimisticUpdate-Exclusive</code><br>(Note 3)         | pessimistic         | update      | serializable  |
| <code>wsPessimisticUpdate-NoCollision</code><br>(Note 4)       | pessimistic         | update      | read committed  |
| <code>wsPessimisticUpdate-WeakestLockAtLoad</code><br>(Note 5) | pessimistic         | update      | Repeatable read   |
| <code>wsOptimisticRead</code>                                  | optimistic          | read        | read committed  |
| <code>wsOptimisticUpdate</code><br>(Note 6)                    | optimistic          | update      | read committed  |

| Profile name  | Concurrency control | Access type | Transaction isolation |
|---|---------------------|-------------|-----------------------|
| <p><b>Notes:</b></p> <ol style="list-style-type: none"> <li>1. Read locks are held for the duration of the transaction.</li> <li>2. The generated SELECT FOR UPDATE query grabs locks at the beginning of the transaction.</li> <li>3. SELECT FOR UPDATE is generated; locks are held for the duration of the transaction.</li> <li>4. <ul style="list-style-type: none"> <li><b>5.0.2</b></li> </ul> <p>A plain SELECT query is generated. No locks are held, but updates are permitted. Use cautiously. This intent enables execution without concurrency control.</p> </li> <li>5. <ul style="list-style-type: none"> <li><b>5.0.2</b></li> </ul> <p>Where supported by the backend, the generated SELECT query does not include FOR UPDATE; locks are escalated by the persistent store at storage time if updates were made. Otherwise, the same as wsPessimisticUpdate.</p> </li> <li>6. Generated overqualified-update query forces failure if CMP column values have changed since the beginning of the transaction. <ul style="list-style-type: none"> <li><b>5.0.2</b></li> </ul> <p>Be sure to review the rules for forming overqualified-update query predicates. Certain column types (for example, BLOB) are ineligible for inclusion in the overqualified-update query predicate and might affect your design.</p> </li> </ol> |                     |             |                       |

## Access intent best practices

This topic outlines issues to consider when applying access intent policies to Enterprise JavaBeans (EJB) methods.

- 

### 5.0.2

**Start by configuring the default access intent policy for an entity.** After your application is built and running, you can more finely tune certain access paths in your application using application profiling or method-level access intent.

- 

### 5.0.2

**Don't mix access types.** Avoid using both pessimistic and optimistic policies in the same transaction. For most databases, pessimistic and optimistic policies use different isolation levels. This can result in multiple database connections, which prevents you from taking advantage of the performance benefits possible through connection sharing.

- **Take care when applying wsPessimisticUpdate-NoCollision.** This policy does not ensure data integrity. No database locks are held, so concurrent transactions can overwrite each other's updates. Use this policy only if you can be sure that only one transaction will attempt to update persistent store at any given time.

## Frequently asked questions: Access intent

I have not applied any access intent policies at all. My application runs just fine with a DB2 database, but it fails with an Oracle database with the following message: `com.ibm.ws.ejbpersistence.utilpm.PersistenceManagerException: PMGR1001E: No such DataAccessSpec :FindAllCustomers. The backend datastore`

**does not support the SQLStatement needed by this AccessIntent: (pessimistic update-weakestLockAtLoad)(collections: transaction/25) (resource manager prefetch: 0) (AccessIntentImpl@d23690a). Why?**

If you have not configured access intent, all of your data is accessed under the default access intent policy (wsPessimisticUpdate-WeakestLockAtLoad). On DB2 databases, the weakest lock is a shared one, and the query runs without a FOR UPDATE clause. On Oracle databases, however, the weakest lock is an update lock; this means that the SQL query must contain a FOR UPDATE clause. However, not every SQL statement necessarily supports FOR UPDATE; for example, if the query is being run against multiple tables in a join, FOR UPDATE is not supported.

To avoid this problem, try either of the following:

- Modify your SQL query or reconfigure your application so that an update lock is supported
- Apply an access intent policy that supports optimistic concurrency

**I am calling a finder method and I get an InconsistentAccessIntentException at run time. Why?**

### 5.0.2

This can occur when you use method-level access intent policies to apply more control over how a bean instance is loaded. This exception indicates that the entity bean was previously loaded in the same transaction. This could happen if you called a multifinder method that returned the bean instance with access intent policy X applied; you are now trying to load the second bean again by calling its `findByPrimaryKey` method with access intent Y applied. Both methods must have the same access intent policy applied.

Likewise, if the entity was loaded once in the transaction using an access intent policy configured on a finder, you might have called a container-managed relationship (CMR) accessor method that returned the entity bean configured to load using that entity's default access intent.

To avoid this problem, ensure that your code does not load the same bean instance twice within the same transaction with different access intent policies applied. Avoid the use of method-level access intent unless absolutely necessary.

**I have two beans in a container-managed relationship. I call `findByPrimaryKey()` on the first bean and then call `getBean2()`, a CMR accessor method, on the returned instance. At that point, I get an `InconsistentAccessIntentException`.**

**Why?** You are probably using read-ahead. When you loaded the first bean, you caused the second bean to be loaded under the access intent policy applied to the finder method for the first bean. However, you have configured your CMR accessor method from the first bean to the second with a different access intent policy. CMR accessor methods are really finder methods in disguise; the run-time environment behaves as if you were trying to change the access intent for an instance you have already read from persistent store.

To avoid this problem, beans configured in a read-ahead hint are all driven to load with the same access intent policy as the bean to which the read-ahead hint is applied.

**I have a bean with a one-to-many relationship to a second bean. The first bean has a pessimistic-update intent policy applied. When I try to add an instance of the second bean to the first bean's collection, I get an**

### **UpdateCannotProceedWithIntegrityException. Why?**

The second bean probably has a read intent policy applied. When you add the second bean to the first bean's collection, you are not updating the first bean's state, you are implicitly modifying the second bean's state. (The second bean contains a foreign key to the first bean, which is modified.)

To avoid this problem, ensure that both ends of the relationship have an update intent policy applied if you expect to change the relationship at run time.

---

## **EJB modules**

An EJB module is used to assemble one or more enterprise beans into a single deployable unit. An EJB module is stored in a standard Java archive (JAR) file.

An EJB module contains the following:

- One or more deployable enterprise beans.
- A deployment descriptor, stored in an Extensible Markup Language (XML) file. This file declares the contents of the module, defines the structure and external dependencies of the beans in the module, and describes how the beans are to be used at run time.

An EJB module can be used as a stand-alone application, or it can be combined with other EJB modules, or with Web modules, to create a J2EE application. An EJB module is installed and run in an enterprise bean container.

For more information about EJB modules, see "Resources for learning."

---

## **Assembling EJB modules**

### Before you begin

If you want to use existing Java 2 Platform, Enterprise Edition (J2EE) Version 1.2 modules in your J2EE Version 1.3 application, migrate them to the Version 1.3 specification first.

Assemble an Enterprise JavaBeans (EJB) module to contain enterprise beans and related code artifacts. Group Web components, client code, and resource adapter code in separate modules.

An EJB module is installed as a stand-alone application or is combined with other modules into an enterprise application.

To increase performance, break CMP enterprise beans into several enterprise bean modules during assembly. The load time for hundreds of beans is improved by distributing the beans across several JAR files and packaging them to an EAR file. Load time is faster when the administrative server attempts to start the beans, for example, 8-10 minutes versus more than one hour when one JAR file is used.

The Application Assembly Tool (AAT) provides flexibility in assembling EJB modules. Options described below include:

- Importing an existing EJB module (EJB JAR file)
- Creating a new EJB module
- Copying code artifacts (such as entity beans) from one EJB module into a new EJB module

### Steps for this task

1. (Start the AAT).
2. From the New tab, select **EJB Module**. Click **OK**.  
The navigation tree displays various sets of properties for configuring the new EJB module.
3. **(Optional)** Use the property dialog shown in the AAT workspace to change the default file name and location.
  - a. It is recommended that you change the display name so that it differs from the file name.
  - b. If you like, change the temporary location of the EJB module from the default location, *install\_root/bin*.
4. Add at least one EJB component to the module.
  - Add at least one enterprise bean to the EJB component.
    - Import an existing JAR or EAR file containing EJB components.
      - a. In the Navigation pane, right-click the **EJB Components** icon.
      - b. Select **Import** from the pop-up menu.
      - c. Click **Browse** to locate the archive file to import.
      - d. Click **Open** to display the contents of the archive file. The applications in the selected archive file display.
      - e. Select an EJB application from the archive file.
      - f. Select the servlets or JSP files to be added, and click **Add** to display the components in the Selected Components window.
      - g. Click **OK** to add the selected components.
    - Copy and paste values from an existing module.
    - Create a new EJB component.
      - a. In the Navigation pane, right-click the **EJB Components** icon.
      - b. Select **New** from the pop-up menu.
      - c. Enter the component name and archive type.
      - d. Select the class files.
      - e. Click **OK** in the **New EJB Component** property dialog.
      - f. Enter properties for the EJB component as needed.
5. Enter assembly properties for each bean.
  - a. Click the plus sign (+) next to the component instance to show property groups.
  - b. Right-click the icon for a property group.
  - c. Select **New** from the pop-up menu to add new values, or edit existing values in the property pane.

If you change the session type (stateful or stateless) of a session bean, you must click **Apply** before making any other changes to the same bean. Otherwise, certain input fields on the GUI become inactive. (You will know they are inactive because they are grayed out on the GUI.) After making all of your changes, click **Apply** again to commit them.

6. Add any other files needed by the application.
  - a. Right-click the **Files** icon.
  - b. Select **Add Files** from the pop-up menu.
  - c. Select **Browse** to navigate the directory structure.
  - d. Click **Select** to open an archive.

- e. Select the files to add and click **Add**.
- f. In the Selected Files window, click **OK** to add the files.

### What to do next

Assemble any other new modules of your choice:

- EJB modules
- (Application client modules)
- Resource adapter modules

You can also (migrate existing modules).

Another option is to proceed directly to (assembling a new application module). While assembling an application module, you can create any new modules that you need.

## **CMP field assembly settings**

In Enterprise JavaBeans (EJB) Version 1.1-compliant beans, container-managed persistence (CMP) fields define the variables in the bean class for which the container must handle persistence management. In EJB Version 2.0-compliant beans, these are replaced by abstract get and set methods; generated code provides the implementation of these abstract methods.

### **Name**

Specifies a subset of public variables in the enterprise bean's implementation class.

## **Container transactions**

Container transaction properties specify how an EJB container is to manage transaction scopes for the enterprise bean's method invocations. A transaction attribute is mapped to one or more methods.

## **Container transaction assembly settings**

Container transaction settings specify how an EJB container is to manage transaction scopes for the enterprise bean's method invocations. Specify one or more methods and associate a transaction attribute with each method.

### **Name**

Specifies a name for the mapping between a transaction attribute and one or more methods.

### **Description**

Contains text that describes the mapping.

### **Transaction attribute**

Specifies how the container must manage the transaction boundaries when delegating a method invocation to an enterprise bean's business method.

|           |  |
|-----------|--|
| Data type | String   |
| Default   | Required   |
| Range     | For all but message-driven beans, valid values are Mandatory, Never, Not Supported, Required, Requires New, Supports. For session beans, Bean Managed is also valid. For message-driven beans, only Bean Managed, Not Supported, and Required are valid. |



Additional information about valid values follows:

### **Bean Managed**

Notifies the container that the bean class directly handles transaction demarcation. This setting can be specified for session beans and (in EJB 2.0 implementations only) for message-driven beans, and it cannot be specified for individual bean methods.

### **Mandatory**

Directs the container to always call the bean method within the transaction context associated with the client. If the client attempts to invoke the bean method without a transaction context, the container throws the `javax.jts.TransactionRequiredException` exception to the client. The transaction context is passed to any EJB object or resource accessed by an enterprise bean method.

EJB clients that access these entity beans must do so within an existing transaction. For other enterprise beans, the enterprise bean or bean method must implement the Bean Managed value or use the Required or Requires New value. For non-enterprise bean EJB clients, the client must access a transaction by using the `javax.transaction.UserTransaction` interface.

**Never** Directs the container to invoke bean methods without a transaction context.

- If the client calls a bean method from within a transaction context, the container throws the `java.rmi.RemoteException` exception.
- If the client calls a bean method from outside a transaction context, the container behaves in the same way as if the Not Supported transaction attribute was set. The client must call the method without a transaction context.

### **Not Supported**

Directs the container to call the bean method without a transaction context. If a client calls a bean method from within a transaction context, the container suspends the association between the transaction and the current thread before invoking the method on the enterprise bean instance. The container then resumes the suspended association when the method invocation returns. The suspended transaction context is not passed to any enterprise bean objects or resources that are used by this bean method.

### **Required**

Directs the container to call the bean method within a transaction context. If a client calls a bean method from within a transaction context, the container calls the bean method within the client transaction context. If a client calls a bean method outside a transaction context, the container creates a new transaction context and calls the bean method from within that context. The transaction context is passed to any enterprise bean objects or resources that are used by this bean method.

### **Requires New**

Directs the container to always call the bean method within a new transaction context, regardless of whether the client calls the method within or outside a transaction context. The transaction context is passed to any enterprise bean objects or resources that are used by this bean method.

### **Supports**

Directs the container to call the bean method within a transaction context if the client calls the bean method within a transaction. If the client calls the

bean method without a transaction context, the container calls the bean method without a transaction context. The transaction context is passed to any enterprise bean objects or resources that are used by this bean method.

### **Methods - Name**

Specifies the name of an enterprise bean method, or the asterisk character (\*). The asterisk is used to denote all methods of an enterprise bean's remote and home interfaces.

### **Methods - Enterprise bean**

Specifies which enterprise bean contains the methods indicated in the **Name** setting.

### **Methods - Type**

Used to distinguish between a method with the same signature that is defined in both the home and remote interface. Use `Unspecified` if a transaction attribute applies to all methods of the bean.

Data type  
Range

String  
Valid values for EJB 1.1 implementations are `Home`, `Remote`, or `Unspecified`. For EJB 2.0 implementations, `Local` and `LocalHome` are also valid.

### **Methods - Parameters**

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

## **EJB module assembly settings**

An EJB module is used to assemble enterprise beans into a single deployable unit. An EJB module contains one or more enterprise beans and a deployment descriptor.

### **File name**

Specifies the file name of the EJB module, relative to the top level of the application package.

### **Alternate DD**

Specifies a deployment descriptor to be used at run time instead of the one installed in the module.

### **Classpath**

The path that contains additional classes required by the application that are not contained in the module's archive file. The class loader uses this path. Specify the values relative to the root of the EAR file and separate the values with spaces. Absolute values that refer to files or directories on the hard drive are ignored.

To specify classes that are not in JAR files but are in the root of the EAR file, use a period and forward slash (`./`). Consider the following example directory structure in which the file `myapp.ear` contains an EJB module named `myejb.jar`. Additional classes reside in `class1.jar` and `class2.zip`. A class named `xyz.class` is not packaged in a JAR file but is in the root of the EAR file.

```
myapp.ear/myejb.jar
myapp.ear/class1.jar
myapp.ear/class2.zip
myapp.ear/xyz.class
```

Specify `class1.jar class2.zip ./` as the value of the Classpath setting. (Name only the directory for .class files.)

### **Display name**

Specifies a short name that is intended to be displayed by GUIs.

### **Description**

Contains text that describes the module.

### **EJB client JAR**

Specifies the location of a JAR file that contains a subset of deployed classes needed by the client.

### **Small icon**

Specifies the name of a JPEG or GIF file that contains a small image (16x16 pixels). The image is used as an icon to represent the module in a GUI.

### **Large icon**

Specifies the name of a JPEG or GIF file that contains a large image (32x32 pixels). The image is used as an icon to represent the module in a GUI.

### **Generalizations - Subtype**

Information about this property is not available.

This property is an IBM extension to the standard J2EE deployment descriptor.

### **Generalizations - Supertype**

Information about this property is not available.

This property is an IBM extension to the standard J2EE deployment descriptor.

### **EJB relationships - Name**

The logical name for a container-managed relationship between EJB 2.0-compliant entity beans.

### **Default data source - JNDI name**

Specifies the default JNDI name for the data source. This default is used if binding information is not specified in the deployment descriptor for an individual enterprise bean.

### **Default CMP connection factory**

Specifies the JNDI name for a CMP connection factory. This setting is applicable only for EJB 2.x-compliant CMP beans.

### **Default authorization - User ID**

Specifies the default user ID for connecting to an enterprise bean's data store.

### **Default authorization - Password**

Specifies the default password for connecting to an enterprise bean's data store.

## **Entity bean assembly settings**

An entity bean encapsulates persistent data, which is stored in a data source, and associated methods to manipulate that data.

### **EJB name**

Specifies a logical name for the enterprise bean. This name must be unique within the EJB module. There is no relationship between this name and the JNDI name.

**Display name**

Specifies a short name that is intended to be displayed by GUIs.

**Description**

Contains text that describes the entity bean.

**EJB class**

Specifies the full name of the enterprise bean class (for example, `com.ibm.ejs.doc.account.AccountBean`).

**Remote - Home**

(Required for EJB 1.x) Specifies the full name of the enterprise bean's home interface class (for example, `com.ibm.ejs.doc.account.AccountHome`).

**Remote - Interface**

(Required for EJB 1.x) Specifies the full name of the enterprise bean's remote interface class (for example, `com.ibm.ejs.doc.account.Account`).

**Local interface - Home**

(Required for EJB 1.x) Specifies the full name of the enterprise bean's local home interface class (for example, `com.ibm.ejs.doc.account.AccountLocalHome`).

**Local interface - Interface**

(Required for EJB 1.x) Specifies the full name of the enterprise bean's local interface class (for example, `com.ibm.ejs.doc.account.AccountLocal`).

**Persistence type**

Specifies whether an entity bean manages its own persistent storage or whether storage is managed by the container.

|           |  |
|-----------|--|
| Data type | String   |
| Range     | Valid values are Bean managed and Container managed. |

**Reentrant**

Specifies whether the entity bean is reentrant. If an enterprise bean is reentrant, it can call methods on itself or call another bean that calls a method on the calling bean. Only entity beans can be reentrant.

If an entity bean is not reentrant and a bean instance is executing a client request in a transaction context and another client using the same transaction context makes a request on the same bean instance, the EJB container throws the `java.rmi.RemoteException` exception to the second client. If a bean is reentrant, the container cannot distinguish this type of illegal loopback call from a legal concurrent call, so the bean must be coded to detect illegal loopback calls.

**Primary key class**

Specifies the full name of the bean's primary key class (for example, `com.ibm.ejs.doc.account.AccountKey`). Composite primary keys map to multiple fields in the entity bean class (or to data structures built from the primitive Java data types) and must be encapsulated in a primary key class.

More complicated enterprise beans are likely to have composite primary keys, with multiple instance variables representing the primary key. A subset of the container-managed fields is used to define the primary key class associated with each instance of an enterprise bean.

### Primary key field

Specifies the name of a simple primary key. Simple primary keys map to a single field in the entity bean class and are made up of primitive Java data types (such as integer or long). If exactly one CMP field is the primary key, it can be specified here.

|           |   |
|-----------|---|
| Data type | String  |
| Range     | Valid values are the name of any one CMP field or Compound key, which appears when the primary key class is set |

### Version

Specifies the version of EJB specification with which a container-managed persistence (CMP) entity bean complies.

|           |                             |
|-----------|-----------------------------|
| Data type | String                      |
| Range     | Valid values are 1.x or 2.x |

### Abstract schema name

Specifies the name of the abstract schema type of an EJB Version 2.x CMP entity bean. It is used in EJB Query Language (QL) queries.

For example, the abstract schema name might be Order for an entity bean whose local interface is com.acme.commerce.Order.

### Small icon

Specifies the name of a JPEG or GIF file that contains a small image (16x16 pixels). The image is used as an icon to represent the entity bean in a GUI.

### Large icon

Specifies the name of a JPEG or GIF file that contains a large image (32x32 pixels). The image is used as an icon to represent the entity bean in a GUI.

### Security identity

Specifies that a principal's credential properties are to be handled as indicated in the Run-As mode property. If this setting is enabled, the Run-As mode property can be edited.

### Run-As mode

Specifies the credential information to be used by the security service to determine the permissions that a principal has on various resources.

At appropriate points, the security service determines whether the principal is authorized to use a particular resource based on the principal's permissions. If the method call is authorized, the security service acts on the principal's credential properties according to the Run-As mode setting of the enterprise bean.

|           |   |
|-----------|---|
| Data type | Enumerated integer  |
| Range     | Valid values are Use identity of caller and Use identity assigned to specified role |

Additional information about valid settings follows:

#### Use identity of caller

The security service makes no changes to the principal's credential properties.

### Use identity assigned to specified role

A principal that has been assigned to the specified security role is used for the execution of the bean's methods. This association is part of the application binding in which the role is associated with a user ID and password of a user who is granted that role.

### Role name

Specifies the name of a security role. If **Run-As mode** is set to Use identity assigned to specified role, a principal that has been granted this role is used.

### Description

Contains further information about the security role.

### Concurrency control

Specifies how the bean is to handle concurrent access to its data. This setting is applicable only for EJB 1.x-compliant entity beans.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |  |
|-----------|--|
| Data type | String                                     |
| Range     | Valid values are Optimistic or Pessimistic |

### Inheritance root

Specifies whether the enterprise bean is at the root of an inheritance hierarchy.

This property is an IBM extension to the standard J2EE deployment descriptor.

### Bean Cache - Activate at

Specifies the point at which an enterprise bean is activated and placed in the cache. Removal from the cache and passivation is also governed by this setting.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |  |
|-----------|--|
| Data type | String   |
| Default   | Transaction  |
| Range     | Valid values are Once, Transaction, and Activity session |

More information about valid values follows:

**Once** Indicates that the bean activates when it is first accessed in the server process, and passivates (and is removed from the cache) at the discretion of the container, for example, when the cache becomes full.

#### Transaction

Indicates that the bean activates at the start of a transaction and passivates (and is removed from the cache) at the end of the transaction.

#### Activity session

Indicates that the bean activates and passivates as follows:

- On an ActivitySession boundary, if an ActivitySession context is present on activation
- On a transaction boundary, if a transaction context (but no ActivitySession context) is present on activation
- Otherwise, on an invocation boundary

The values of the **Activate at** and **Load at** settings govern which commit options are used, as follows:

- For commit option A (implies exclusive database access), use `Activate at = Once` and `Load at = Activation`.

This option reduces database I/O (avoids calls to the `ejbLoad` function) but serializes all transactions accessing the bean instance. Option A can increase memory usage by maintaining more objects in the cache, but can provide better response time if bean instances are not generally accessed concurrently by multiple transactions. To use Option A successfully, you must also set Concurrency control to `Pessimistic`.

**Note for Network Deployment users:** When workload management is enabled, you cannot use Option A. You must use settings that result in the use of options B or C.

- For commit option B (implies shared database access), use `Activate at = Once` and `Load at = Transaction`.

Option B can increase memory usage by maintaining more objects in the cache. However, because each transaction creates its own copy of an object, there can be multiple copies of an instance in memory at any given time (one per transaction), requiring database access at each transaction. If an enterprise bean contains a significant number of calls to the `ejbActivate` function, using Option B is beneficial because the required object is already in the cache. Otherwise, this option does not provide significant benefits over Option A.

- For commit option C (implies shared database access), use `Activate at = Transaction` and `Load at = Transaction`.

This option reduces memory usage by maintaining fewer objects in the cache; however, there can be multiple copies of an instance in memory at any given time (one per transaction). This option can reduce transaction contention for enterprise bean instances that are accessed concurrently but not updated.

### Bean Cache - Load at

Specifies when the bean loads its state from the database. The value of this setting implies whether the container has exclusive or shared access to the database.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |   |
|-----------|---|
| Data type | String  |
| Default   | Transaction   |
| Range     | Valid values are <code>Activation</code> and <code>Transaction</code> |

Additional information about valid values follows:

#### Activation

Indicates that the bean loads when it is activated (regardless of `Activate at` setting) and implies that the container has exclusive access to the database.

#### Transaction

Indicates that the bean loads at the start of a transaction and implies that the container has shared access to the database.

The `Activate at` and `Load at` settings govern which commit options are used. The commit options are described in the Enterprise JavaBeans specification. For more information about this setting and achieving a given commit behavior, see `Bean Cache - Activate at`.

### Commit option

Specifies which commit option is used as a result of bean cache settings. The commit options are described in the Enterprise JavaBeans specification.

|           |                              |
|-----------|------------------------------|
| Data type | String                       |
| Range     | Valid values are A, B, and C |

### Local Transactions - Unresolved action

Specifies the action that the EJB container must take if resources are uncommitted by an application in a local transaction.

This property is an IBM extension to the standard J2EE deployment descriptor. This setting is applicable only when **Resolution control** is set to `Application`. A local transaction context is created when a method runs in what the EJB specification refers to as an unspecified transaction context.

|           |  |
|-----------|--|
| Data type | String   |
| Default   | Rollback   |
| Range     | Valid values are <code>Commit</code> and <code>Rollback</code> |

Additional information about these settings follows:

#### Commit

At end of the local transaction context, the container instructs all unresolved local transactions to commit.

#### Rollback

(Default) At end of the local transaction context, the container instructs all unresolved local transactions to roll back.

### Local Transactions - Resolution control

Specifies how the local transaction is to be resolved before the local transaction context ends: by the application through user code or by the EJB container.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |  |
|-----------|--|
| Data type | String   |
| Range     | Valid values are <code>Application</code> and <code>ContainerAtBoundary</code> |

Additional information about these settings follows:

#### Application

When this setting is used, your code must either commit or roll back the local transaction. If this does not occur, the runtime environment logs a warning and automatically commits or rolls back the connection as specified by the **Unresolved action** setting.

#### ContainerAtBoundary

When this setting is used, the container takes responsibility for resolving each local transaction. This provides you with a programming model similar to global transactions in which your code simply gets a connection and performs work within it. User code does not have to handle local transactions.

- If the **Boundary** attribute is set to `ActivitySession`, then the local transactions are enlisted as `ActivitySession` resources and directed to complete by the `ActivitySession`.
- If the the **Boundary** attribute is set to `BeanMethod`, then the local transactions are committed at method end by the container.



Connections are never committed automatically by the resource adapter when this value is configured for the bean **Unresolved action** is not used. An application cannot call `Connection.LocalTransaction.begin()` when using this policy and receives an exception from the resource adapter if it does so.

When using a **Resolution control** of `ContainerAtBoundary`, applications must get connection handles *after* the local transaction context boundary has been started by the container. The application should close the connection before the end of the boundary, although any work performed on the connection is not committed or rolled back until the local transaction context ends. This model of connection usage is sometimes referred to as the "get-use-close" model.

This value is supported only for EJB components that use container-managed transactions. It is not supported for web components or for enterprise beans that use bean-managed transactions.

### **Local Transactions - Boundary**

Specifies the duration of a local transaction context.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |   |
|-----------|---|
| Data type | String  |
| Default   | BeanMethod                                      |
| Range     | Valid values are BeanMethod and ActivitySession |

Additional information about valid settings follows:

#### **BeanMethod**

When this setting is used, the local transaction begins when the method begins and ends when the method ends.

#### **ActivitySession**

When this setting is used, the local transaction must be resolved within the scope of any `ActivitySession` in which it was started or, if no `ActivitySession` context is present, within the same bean method in which it was started.

### **Local Relationship Roles - Name**

Within a local relationship between EJB 1.x-compliant entity beans, the logical name for the view an entity bean presents to other beans in the relationship.

For example, in a relationship between `Account` and `Customer` beans, the role of the `Account` instance relative to the `Customer` instance might be *savingsAccount*.

This property is an IBM extension to the standard J2EE deployment descriptor. This is separate from the container-managed relationships defined in the Enterprise JavaBeans specification, Version 2.0.

### **Local Relationship Roles - Source EJB Name**

The name of the entity bean for which the role is defined.

This property is an IBM extension to the standard J2EE deployment descriptor.

### Local Relationship Roles - is Forward

Specifies how deployment code for navigating the relationship is generated. This setting is applicable only for navigable relationships.

If **isForward** is enabled (set to true), deployment code is generated in the source bean. That is, navigation of the relationship proceeds *forward* from the source to the target.

Otherwise, deployment code is generated in the target bean. That is, navigation of the relationship proceeds from the target to the source.

This property is an IBM extension to the standard J2EE deployment descriptor. For more information, see the documentation for the Deployment Tool for Enterprise JavaBeans.

### Local Relationship Roles - is Navigable

Specifies whether data in related beans may be retrieved through queries to the source bean.

This property is an IBM extension to the standard J2EE deployment descriptor.

### Lifetime in cache

The lifetime, in seconds, of cached data for an instance of this bean type.

This value indicates how long the cached data is to exist beyond the end of the transaction in which the data was retrieved. This might avoid another retrieval from persistent storage if the same bean instance were to be used in later transactions. How this value is interpreted depends on the value of **Lifetime in cache usage**.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |                   |
|-----------|-------------------|
| Data type | Long              |
| Units     | Seconds           |
| Default   | 0                 |
| Range     | 0 to $2^{61} - 1$ |

### Lifetime in cache usage

Indicates how the lifetime-in-cache setting is to be used by the caching mechanism.

This property is an IBM extension to the standard J2EE deployment descriptor.

If your application uses CMP beans in which the underlying data changes infrequently, you might gain significantly better performance by using this setting with **Lifetime in cache**. Typically, data read from persistent storage is held temporarily in an internal cache until the state of the instance is restored. Cached data normally does not persist beyond state restoration or the end of the transaction in which the finder method was called. By setting **Lifetime in cache usage** to a value other than 0ff, you indicate that the cached data is to be held for a longer time, potentially hours or days, before invalidating the version of the data in the cache and fetching a new version. Avoiding a trip to persistent storage greatly speeds up access to such beans by applications.

In addition, the use of a value other than `Off` requires that finders on the bean have an access type of `Optimistic Read` (if you are only reading instances of the bean) or `Optimistic Update` (if you plan to occasionally update instances of the bean).

- For EJB 1.x-compliant beans, see `Access intent - access type`.
- For EJB 2.x-compliant beans, see `Applied access intent`.

`Setting Bean Cache - Activate at to activation` and `Bean Cache - Load at to Once` also minimizes retrievals from persistent storage. However, this settings combination might not be supported by certain CMP beans because it results in the `ejbLoad()` method being called once instead of at the beginning of each transaction in which they are used. The `lifetime-in-cache` settings combination is independent of CMP bean implementation, though it does incur the modest overhead of calling `ejbLoad()` on each use.

|           |   |
|-----------|---|
| Data type | Enumerated int  |
| Units     | Not applicable  |
| Default   | 0 ( <code>Off</code> )  |
| Range     | Valid values are <code>Clock Time</code> , <code>Elapsed Time</code> , <code>Week Time</code> , or <code>Off</code> |

Additional information about valid values follows:

**Off** When this value is used, the value of **Lifetime in cache** is ignored. Beans of this type are cached only in a transaction-scoped cache. The cached data for this instance expires after the transaction in which it was retrieved is completed.

#### **Elapsed Time**

When this value is used, the value of **Lifetime in cache** is added to the time at which the transaction in which the bean instance was retrieved is completed. The resulting value becomes the time at which the cached data expires. The value of **Lifetime in cache** can add up to minutes, hours, days, and so on.

#### **Clock Time**

When this value is used, the value of **Lifetime in cache** represents a particular time of day. The value is added to the immediately preceding or following midnight to calculate a future time value, which is then treated as for `Elapsed Time`. Using `Clock Time` enables you to specify that all instances of this bean type are to have their cached data invalidated at, for example, 3 AM, no matter when they were retrieved. This is important if, for example, the data underlying this bean type is batch-updated at 3 AM every day.

The selection of midnight (preceding or following) depends on the value of **Lifetime in cache**. If **Lifetime in cache** plus the value that represents the preceding midnight is earlier than the current time, the following midnight is used.

When you use `Clock Time`, the value of **Lifetime in cache** is not supposed to represent more than 24 hours. If it does, the cache manager subtracts 24-hour increments from it until a value less than or equal to 24 hours is achieved. To invalidate data at midnight, set **Lifetime in cache** to 0.

#### **Week Time**

Usage of this value is the same as for `Clock Time`, except that the value of **Lifetime in cache** is added to the preceding or following Sunday

midnight (11:59 PM Saturday plus 1 minute). When Week Time is used, the value of **Lifetime in cache** can represent more than 24 hours but not more than 7 days.

### **Default Access Intent**

Specifies the default access intent under which the entity should load.

#### **Applicability of the following table: [Fix Pack 5.0.2 and later]**

|           |  |
|-----------|--|
| Data type | String   |
| Units     | Not applicable   |
| Default   | Not applicable   |
| Range     | Valid settings are wsPessimisticUpdate, wsPessimisticUpdate-NoCollision, wsPessimisticUpdate-Exclusive, wsPessimisticUpdate-WeakestLockAtLoad, wsPessimisticRead, wsOptimisticUpdate, or wsOptimisticRead. |

### **JNDI name**

Specifies the JNDI name of the bean's home interface. This is the name under which the enterprise bean's home interface is registered and therefore, is the name that must be specified when an EJB client does a lookup of the home interface.

### **Data source - JNDI name**

Specifies the JNDI name for the bean's data source.

### **Default Authorization - User ID**

Specifies the default user ID for connecting to a data source.

### **Default Authorization - Password**

Specifies the default password for connecting to a data source.

### **CMP Resource - JNDI name**

Specifies the JNDI name for the resource by which CMP data is stored.

### **CMP Resource - Resource authentication**

Specifies the scope at which resources are to be authenticated: by the container or by the resource.

### **WAS Enterprise assembly settings for entity beans**

Use this page to configure Enterprise Application Server functions for entity beans.

**Note:** For entity beans, the internationalization type is always **Container**; it is not configurable.

**Task references:** Task references enable developers to programmatically set the current task under which a component is executing.

**Name:** The name of the task reference. This name corresponds to the string that the developer uses to set the current task.

**Task:** The task that the run-time environment will associate with execution of the current component.

**Name:** The name of the task.

Programmatically setting the logical task name causes this task name to be associated with requests from this J2EE component.

**Description:** A description of the task.

An optional field provided as a convenience for the developer.

## Message-driven bean assembly settings

Use this page to configure the assembly properties of message-driven beans

For more information about the effect of JMS properties, such as message selectors and message acknowledgement, see the WebSphere MQ *Using Java* book, SC34-5456 or Sun's Java Message Service (JMS) specification documentation (<http://developer.java.sun.com/developer/technicalArticles/Networking/messaging/>)

The following notebook pages are available:

### General properties

Specify general assembly properties for the message bean.

- EJB name
- Display name
- Description
- EJB class
- Transaction type

### Advanced properties

Specify advanced assembly properties for the message bean.

- Message selector
- Acknowledge mode
- Destination type

### Bindings properties

Specify bindings assembly properties for the message bean.

- [ ]

### EJB name

The logical name for the message bean (as an enterprise bean)

The logical name for the message bean (as an enterprise bean). This name must be unique within the EJB module. There is no relationship between this name and the JNDI name.

Data type String

### Display name

A short name that is intended to be displayed by graphical user interfaces

Data type String

### Description

A description of the message bean, for administrative use

Data type String

## EJB class

The full package name of the message bean class

Specify the full package name of the message bean class, for example, `com.ibm.ejs.doc.account.MessageBean`. You can either type the class name or click **Browse** to locate an existing class file.

Data type String

## Transaction type

Whether the message bean manages its own transactions or the container manages transactions on behalf of the bean

Whether the message bean manages its own transactions or the container manages transactions on behalf of the bean. All messages retrieved from a specific destination have the same transactional behavior. To enable the transactional behavior that you want, you must configure the JMS destination with the same transactional behavior as you configure for the message bean.

|           |  |
|-----------|--|
| Data type | Enum   |
| Default   | Bean   |
| Range     | <b>Bean</b> The message bean manages its own transactions<br><b>Container</b> The container manages transactions on behalf of the bean |

## Message selector

The JMS message selector to be used to determine which messages the message bean receives

The JMS message selector to be used to determine which messages the message bean receives; for example:

```
JMSType='car' AND color='blue' AND weight>2500
```

The selector string can refer to fields in the JMS message header and fields in the message properties. Message selectors cannot reference message body values.

Data type String  
Range A String whose syntax is based on a subset of the SQL92 conditional syntax.

## Acknowledge mode

How the session acknowledges any messages it receives.

This property applies only to message-driven beans that uses bean-managed transaction demarcation (**Transaction type** is set to Bean).

|           |                  |
|-----------|------------------|
| Data type | Enum             |
| Default   | Auto Acknowledge |

Range

#### **Auto Acknowledge**

The session automatically acknowledges a message when it has either successfully returned from a call to receive, or the message listener it has called to process the message successfully returns.

#### **Dups OK Acknowledge**

The session lazily acknowledges the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

### **Destination type**

Whether the message bean uses a queue or topic destination.

Data type

Enum

Default

Null

Range

**Queue** The message bean uses a queue destination.

**Topic** The message bean uses a topic destination.

### **Listener port name**

The name of the listener port for this message bean.

The name of the listener port for this message bean (as defined on the WebSphere administrative console).

Data type

String

### **WAS Enterprise assembly settings for message-driven beans**

Use this page to configure Enterprise Application Server functions for message-driven beans.

**Task references:** Task references enable developers to programmatically set the current task under which a component is executing.

**Name:** The name of the task reference. This name corresponds to the string that the developer uses to set the current task.

**Task:** The task that the run-time environment will associate with execution of the current component.

**Name:** The name of the task.

Programmatically setting the logical task name causes this task name to be associated with requests from this J2EE component.

**Description:** A description of the task.

An optional field provided as a convenience for the developer.

**Internationalization type:** The internationalization type of the message-driven bean.

The internationalization type of a message-driven bean indicates whether the bean or the hosting EJB container will manage internationalization context on business method invocations. For message-driven beans the internationalization type setting can be configured to **Application** or **Container**, and defaults to **Container**. The setting applies to all business methods of a message-driven bean.

Select whether the message-driven bean will employ Application or Container internationalization context management using the **Internationalization Type** drop-down menu.

**Note:** In this release, messages received by message-driven bean onMessage() invocations lack internationalization context, meaning that caller context elements are unavailable and always default to the elements of the server JVM, regardless of the internationalization type or applicable container internationalization attribute.

Default

Container



Range

## Application

Under Application-managed Internationalization (AMI) a message-driven bean is responsible to manage (set) invocation context using the internationalization context API. The container suspends internationalization context imported from the calling process on business method invocations; to continue propagating this context on subsequent business method requests, a method must use the API to transfer caller context elements to the invocation context scoped to that method. Invocation context elements not set using the API default to the respective elements of the current process when accessed using the API or when propagated on outgoing business method requests.

Specify Application-managed Internationalization for message-driven beans having internationalization context management requirements not supported by Container-managed Internationalization.

## Container

Under Container-managed Internationalization (CMI) a message-driven bean has read-only access to the internationalization context API and relies solely on the container to manage invocation context in accordance with the **Run as** field of the applicable Container Internationalization attribute. On a business method request, the container scopes the context indicated by the **Run as** field to the method; when the method has completed, the container removes this context from scope. If the bean method is not indicated within a container internationalization attribute, the container scopes the caller context by default. Invocation context elements unavailable to the container default to the respective elements of the current JVM when accessed using the API or when propagated on outgoing business method requests.

## Method extensions

Method extensions are IBM extensions to the standard deployment descriptors for enterprise beans.

Method extension properties are used to define transaction isolation levels for methods, to control the delegation of a principal's credentials, and to define custom finder methods.

## Method extension assembly settings

Method extensions are IBM extensions to the standard J2EE deployment descriptors for Enterprise JavaBeans (EJB) Version 1.x-compliant beans. Method extension settings define transaction isolation levels for methods and control the delegation of a principal's credentials.

### Method type

Specifies the type of the enterprise bean method.

|           |   |
|-----------|---|
| Data type | String  |
| Range     | Valid values are Home, Remote, and Unspecified. |

### Name

Specifies the name of an enterprise bean method, or the asterisk character (\*). The asterisk is used to denote all methods of an enterprise bean's remote and home interfaces.

### Parameters

Contains a list of fully qualified Java type names of the method parameters. Used to identify a single method among multiple methods with an overloaded method name.

### Isolation level attributes

The transaction isolation level determines how isolated one transaction is from another. This can be set for individual methods in an enterprise bean or for all methods in the enterprise bean. An asterisk is used to indicate all methods in the bean. This setting is not applicable for EJB 2.x-compliant beans.

Within a transactional context, the isolation level associated with the first method call becomes the required isolation level for all methods called within that transaction. If a method is called with a different isolation level from that of the first method, the `java.rmi.RemoteException` exception is thrown.

### Isolation level

Specifies the level of transactional isolation.

The container uses the transaction isolation level attribute as follows:

- Session beans and entity beans with bean-managed persistence (BMP): For each database connection used by the bean, the container sets the transaction isolation level at the start of each transaction unless the bean explicitly sets the isolation level on the connection.
- Entity beans with container-managed persistence (CMP): The container generates database access code that implements the specified isolation level.

|           |  |
|-----------|--|
| Data type | String   |
| Range     | Valid values are Serializable, Repeatable read, Read committed, and Read uncommitted |

### Serializable

This level prohibits the following types of reads:

- *Dirty* reads, in which a transaction reads a database row containing uncommitted changes from a second transaction.

- *Nonrepeatable* reads, in which one transaction reads a row, a second transaction changes the same row, and the first transaction rereads the row and gets a different value.
- *Phantom* reads, in which one transaction reads all rows that satisfy an SQL WHERE condition, a second transaction inserts a row that also satisfies the WHERE condition, and the first transaction applies the same WHERE condition and gets the row inserted by the second transaction.

#### **Repeatable read**

This level prohibits dirty reads and nonrepeatable reads, but it allows phantom reads.

#### **Read committed**

This level prohibits dirty reads but allows nonrepeatable reads and phantom reads.

#### **Read uncommitted**

This level allows dirty reads, nonrepeatable reads, and phantom reads.

### **Access intent - Intent type**

Specifies whether to load the enterprise bean as read-only or for update. This setting is applicable only for EJB 1.x-compliant beans.

This setting is applicable for the following types of beans:

- EJB 1.x-compliant entity beans
- Enterprise beans with CMP version 1.x that are packaged in EJB 2.x-compliant modules

To specify the access intent for EJB 2.x-compliant beans, select an access intent policy.

|           |                                 |
|-----------|---------------------------------|
| Data type | String                          |
| Range     | Valid values are Read or Update |

### **Finder descriptor - User**

Specifies that the user has provided a finder helper class in the entity bean's home interface. The class contains specialized finder methods. This setting is applicable only for EJB 1.x-compliant entity beans.

### **Finder descriptor - EJB QL**

Describes the semantics of a finder method that uses EJB QL (Enterprise JavaBeans query language). This setting is applicable only for EJB 1.x-compliant entity beans.

EJB QL is a declarative, SQL-like language that is intended to be compiled to the target language of the persistent datastore used by a persistence manager. The language is independent of the bean's mapping to a relational datastore and is therefore portable. The EJB query specifies a search based on the persistent attributes and relationships of the bean. An EJB query can contain the following clauses:

- SELECT (optional), which specifies the EJB objects to return
- FROM (required), which specifies the collections of objects to which the query is to be applied
- WHERE (optional), which contains search predicates over the collections
- ORDER BY (optional), which specifies the ordering of the resulting collection

### **Finder descriptor - Full SELECT**

Describes the semantics of a finder method that uses an SQL SELECT clause. For information on restrictions, see the documentation for the Deployment Tool for Enterprise JavaBeans.

### **Finder descriptor - WHERE clause**

Describes the semantics of a finder method that uses an SQL WHERE clause. This clause restricts the results that are returned by the query. For information on restrictions, see the documentation for the Deployment Tool for Enterprise JavaBeans.

### **Security identity**

Specifies whether a principal's credential settings are to be handled as indicated in the **Run-As mode** setting. If this is enabled, the **Run-As mode** setting can be edited.

### **Description**

Contains further information about the security instructions.

### **Run-As mode**

Specifies the credential information to be used by the security service to determine the permissions that a principal has on various resources.

At appropriate points, the security service determines whether the principal is authorized to use a particular resource based on the principal's permissions. If the method call is authorized, the security service acts on the principal's credential settings according to the **Run-As mode** setting of the enterprise bean.

|           |  |
|-----------|--|
| Data type | Enumerated integer   |
| Range     | Valid values are Use identity of caller, Use identity of EJB server, and Use identity assigned to specified role |

Additional information about valid values for this setting follows:

#### **Use identity of caller**

The security service makes no changes to the principal's credential settings.

#### **Use identity of EJB server**

The security service alters the principal's credential settings to match the credential settings associated with the EJB server.

#### **Use identity assigned to specified role**

A principal that has been assigned to the specified security role is used for the execution of the bean's methods. This association is part of the application binding in which the role is associated with a user ID and password of a user who is granted that role.

### **Role name**

Specifies the name of a security role. If **Run-As mode** is set to Use identity assigned to specified role, a principal that has been granted this role is used.

### **Description**

Contains further information about the security role.

## **Method permissions**

A method permission is a mapping between one or more security roles and one or more methods that a member of the role can call.

## Method permission assembly settings

A method permission is a mapping between one or more security roles and one or more methods that a member of the role can call. Assembly settings for method permissions include an optional description, a list of security role names, and a list of methods. The security roles must be defined, and the methods must be defined in the enterprise bean's remote or home interfaces.

### Method permission name

Specifies a name for the mapping between method permissions and security roles.

### Description

Contains text that describes the mapping between method permissions and security roles.

### Methods - Name

Specifies the name of an enterprise bean method, or the asterisk (\*) character. The asterisk is used to denote all the methods of an enterprise bean's remote and home interfaces.

### Methods - Enterprise bean

Specifies the name of the enterprise bean that contains the method.

### Methods - Type

Distinguishes between a method with the same signature that is defined in both the home and remote interface. Use `Unspecified` if a method permission applies to all methods of a bean.

|           |  |
|-----------|--|
| Data type | String   |
| Range     | Valid values are <code>Unspecified</code> , <code>Remote</code> , or <code>Home</code> . |

### Methods - Parameters

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

### Unchecked

Specifies whether the method permission is checked before the method is run.

### Roles - Role name

Specifies the name of the security role that must be granted in order to call the method.

## Query assembly settings

Use these to specify a finder or SELECT query.

A query element contains the following:

- Optional description of the query
- Name of the finder or SELECT method that uses the query
- The return type of mapping, if it is used
- Whether the query is for a SELECT method
- EJB query language (EJB QL) query string that defines the query

Queries that are expressed in EJB QL must use the `ejb-ql` element to specify the query. If a query cannot be expressed in EJB QL, describe the semantics of the query by using the description element and leave the `ejb-ql` element empty.

**Name**

Contains the name of an enterprise bean method or the asterisk (\*) character. An asterisk in the method-name element denotes all methods of an enterprise bean's remote and home interfaces.

**Parameters**

Contains a list of the fully-qualified Java names of the method parameters.

**Result type**

Used in the query element to indicate whether a returned abstract schema type for a SELECT method should be mapped to an EJBLocalObject or EJBObject type.

**References**

References are logical names used to locate external resources for enterprise applications. References are defined in the application's deployment descriptor file. At deployment, the references are bound to the physical location (global JNDI name) of the resource in the target operational environment.

This product supports the following types of references:

- An EJB reference is a logical name used to locate the home interface of an enterprise bean.
- A resource reference is a logical name used to locate a connection factory object.

These objects define connections to external resources such as databases and messaging systems. The container makes references available in a JNDI naming subcontext. By convention, references are organized as follows:

- EJB references are made available in the `java:comp/env/ejb` subcontext.
- Resource references are made available as follows:
  - JDBC DataSource references are declared in the `java:comp/env/jdbc` subcontext.
  - JMS connection factories are declared in the `java:comp/env/jms` subcontext.
  - JavaMail connection factories are declared in the `java:comp/env/mail` subcontext.
  - URL connection factories are declared in the `java:comp/env/url` subcontext.

**EJB reference assembly settings**

An EJB reference is a logical name used to locate the home interface of an enterprise bean used by an application.

At deployment, the EJB reference is bound to the enterprise bean's home in the target operational environment. The container makes the application's EJB references available in a JNDI naming context. It is recommended that references to enterprise beans be organized in the `ejb` subcontext of the application's environment (in `java:comp/env/ejb`).

**Name**

Specifies the JNDI name of the enterprise bean's home interface relative to the `java:comp/env` context.

For example, if `ejb/EmplRecord` is specified, the referring code looks up the enterprise bean's home interface at `java:comp/env/ejb/EmplRecord`. This JNDI name is an alias used by the code (the actual JNDI name is specified on the Binding tab).

## Description

Contains text that describes the EJB reference.

## Link

Used to link an EJB reference to an enterprise bean in the current module (the same module as the one making the reference) or in another module within the same J2EE application. This setting specifies the name of the target enterprise bean.

The target enterprise bean can be in any EJB module in the same J2EE application as the referring module. To avoid having to rename enterprise beans to have unique names within an J2EE application, specify the path name of the EJB archive file that contains the referenced enterprise bean and append the target bean's name, separated by a # symbol (for example, `...products/product.jar#ProductEJB`). The path name is relative to the referring module's archive file specification. If a link is not specified, the reference must be resolved to a JNDI name during installation.

## Home

Specifies the fully qualified name of the enterprise bean's home interface (for example, `com.ibm.ejbs.EmplRecordHome`).

## Remote

Specifies the fully qualified name of the enterprise bean's remote interface (for example, `com.ibm.ejbs.EmplRecord`).

## Type

Specifies the expected type of the referenced enterprise bean.

|           |                   |
|-----------|-------------------|
| Data type | String            |
| Default   | None; must be set |
| Range     | Entity or Session |

## JNDI name

Binding information that is used by the run-time environment to resolve the location of a resource.

For EJB references, the value of this setting must match the JNDI name of the enterprise bean as it was specified on the **Binding** tab for the EJB module that contains the bean.

## EJB local-reference assembly settings

For EJB 2.0-compliant beans, the EJB local reference element declares a reference to another enterprise bean's local home interface.

## Name

Specifies the name of an EJB reference.

This is the JNDI name that the servlet code uses to get a reference to the enterprise bean. The following example illustrates how this element is specified in the deployment descriptor:

```
<ejb-ref-name>ejb/Payroll</ejb-ref-name>
```

## Description

Contains a description of the parent element.

This can include any information that the EJB archive-file producer wants to provide to the consumer of the EJB archive file.

### **Link**

Used in the `ejb-ref` element to specify that an EJB reference is linked to an enterprise bean in the encompassing web-application package.

The value of the `link` element must be the EJB name of an enterprise bean in the same web-application package. The following example illustrates how this element is specified in the deployment descriptor:

```
<ejb-link>EmployeeRecord</ejb-link>
```

### **Local interface**

Specifies the fully-qualified name of the enterprise bean's local interface.

### **Local home**

Specifies the fully-qualified name of the enterprise bean's local home interface.

### **Type**

Specifies the expected type of the referenced enterprise bean.

## **EJB relation assembly settings**

An EJB relation describes a relationship between two entity beans with container-managed persistence.

The name of the relationship, if specified, is unique within an EJB archive file.

### **Description**

Contains text to describe the EJB relationship role.

### **Source EJB**

Specifies the source of the role that participates in a relationship.

### **Multiplicity**

Specifies the multiplicity of the role that participates in a relation.

### **Cascade delete**

Within a particular relationship, specifies that the lifetime of one or more entity beans is dependent on the lifetime of another entity bean.

**Cascade delete** can be specified only for an EJB relationship role contained in an EJB relation in which the other EJB relationship role specifies a multiplicity of one.

### **CMR field**

Enables the declaration of a container-managed relationship (CMR) field.

The CMR field describes the bean provider's view of a relationship. It consists of an optional description and the name and class type of the source enterprise bean's role in a relationship.

## **Exclude list assembly settings**

The exclude list indicates which methods in the enterprise beans may not be called. You should also configure security for the enterprise bean so that access to the listed methods is not permitted.

This capability applies only to Enterprise JavaBeans (EJB) Version 2.x-compliant beans. For more information about exclude lists, see the EJB specification.



## Description

Provides additional information about this exclude list.

## Methods - Name

Specifies the name of an enterprise bean method, or the asterisk (\*) character. The asterisk is used to denote all the methods of an enterprise bean's remote and home interfaces.

## Methods - Enterprise bean

Specifies the name of the enterprise bean that contains the method.

## Methods - Type

Distinguishes between a method with the same signature that is defined in both the home and remote interface. Use Unspecified if the exclusion applies to all methods of a bean.

|           |   |
|-----------|---|
| Data type | String  |
| Range     | Valid values are Unspecified, Remote, or Home |

## Methods - Parameters

Contains a list of fully qualified Java type names of the method parameters. This setting is used to identify a single method among multiple methods with an overloaded method name.

## Security role assembly settings

A security role is a logical grouping of principals. Access to operations (such as enterprise-bean methods) is controlled by granting access to a role.

### Role name

Specifies the name of a security role that is unique to an application. This setting applies only when you are specifying security roles at the application level (EAR file).

### Description

Contains text that describes the application-specific security role. This setting applies only when you are specifying security roles at the application level (EAR file).

### Binding - Groups - Name

Specifies the user groups that are granted the application-specific security role. This setting applies only when you are specifying security roles at the application level (EAR file).

### Binding - Users - Name

Specifies the users that are granted the application-specific security role. This setting applies only when you are specifying security roles at the application level (EAR file).

### Binding - Special Subjects - Name

Specifies one of two special categories of authenticate users to which application-specific security roles can be granted: Everyone or All. This setting applies only when you are specifying security roles at the application level (EAR file).

If the special subject All is granted a role, any user who can authenticate by using a valid user ID and password is considered to be granted that role.

If the special subject Everyone is granted a role, all users, including those who did not authenticate, are granted the role. In other words, a method on an enterprise bean or a URI is unprotected if any of the required roles for that method are granted to the special subject Everyone.

|           |                                  |
|-----------|----------------------------------|
| Data type | String                           |
| Range     | Valid values are All or Everyone |

## Session bean assembly properties

A session bean encapsulates transient data that is associated with a particular EJB client. Unlike data in an entity bean, the data in a session bean is not stored in a persistent data source.

### EJB name

Specifies a logical name for the enterprise bean. This name must be unique within the EJB module. There is no relationship between this name and the JNDI name.

### Display name

Specifies a short name that is intended to be displayed by GUIs.

### Description

Contains text that describes the session bean.

### EJB class

Specifies the full name of the enterprise bean class (for example, com.ibm.ejs.doc.account.AccountBean).

### Remote - Home

Specifies the full name of the enterprise bean's home interface class (for example, com.ibm.ejs.doc.account.AccountHome).

### Remote - Interface

Specifies the full name of the enterprise bean's remote interface class (for example, com.ibm.ejs.doc.account.Account).

### Local interface - Home

Specifies the full name of the enterprise bean's home interface class (for example, com.ibm.ejs.doc.account.AccountLocalHome).

### Local interface - Interface

Specifies the full name of the enterprise bean's local interface class (for example, com.ibm.ejs.doc.account.AccountLocal).

### Session type

Specifies whether the enterprise bean maintains a conversational state (is stateful) or does not (is stateless).

|           |   |
|-----------|---|
| Data type | String                                  |
| Range     | Valid values are Stateful and Stateless |

### Transaction type

Specifies whether the enterprise bean manages its own transactions or whether the container manages transactions on behalf of the bean.

|           |                                    |
|-----------|------------------------------------|
| Data type | String                             |
| Range     | Valid values are Container or Bean |

### Small icon

Specifies the name of a JPEG or GIF file that contains a small image (16x16 pixels). The image is used as an icon to represent the session bean in a GUI.

### Large icon

Specifies the name of a JPEG or GIF file that contains a large image (32x32 pixels). The image is used as an icon to represent the session bean in a GUI.

### Security identity

Specifies whether a principal's credential properties are to be handled as indicated in the **Run-As mode** property. If this setting is enabled (that is, set to true), the **Run-As mode** setting can be edited.

### Description

Contains further information about the security instructions.

### Run-As mode

Specifies the credential information to be used by the security service to determine the permissions that a principal has on various resources.

At appropriate points, the security service determines whether the principal is authorized to use a particular resource based on the principal's permissions. If the method call is authorized, the security service acts on the principal's credential properties according to the **Run-As mode** setting of the enterprise bean.

|           |   |
|-----------|---|
| Data type | Enumerated integer  |
| Range     | Valid values are Use identity of caller and Use identity assigned to specified role |

Additional information about valid values for this setting follows:

#### Use identity of caller

The security service makes no changes to the principal's credential properties.

#### Use identity assigned to specified role

A principal that has been assigned to the specified security role is used for the execution of the bean's methods. This association is part of the application binding in which the role is associated with a user ID and password of a user who is granted that role.

### Role name

Specifies the name of a security role. If **Run-As mode** is set to Use identity assigned to specified role, a principal that has been granted this role is used.

### Description

Contains further information about the security role.

### Timeout

This property applies only to stateful session beans.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |         |
|-----------|---------|
| Data type | Integer |
| Units     | Seconds |

### **Inheritance root**

Specifies whether the enterprise bean is at the root of an inheritance hierarchy.

This property is an IBM extension to the standard J2EE deployment descriptor.

### **Bean Cache - Activate at**

Specifies the point at which an enterprise bean is activated and placed in the cache. Removal from the cache and passivation is also governed by this setting. This setting applies to stateful session beans only (not to stateless beans).

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |                                       |
|-----------|---------------------------------------|
| Data type | String                                |
| Default   | Once                                  |
| Range     | Valid values are Once and Transaction |

Additional information about valid values follows:

**Once** Indicates that the bean is activated when it is first accessed in the server process, and passivated (and removed from the cache) at the discretion of the container, for example, when the cache becomes full.

#### **Transaction**

Indicates that the bean is activated at the start of a transaction and passivated (and removed from the cache) at the end of the transaction.

### **Local Transactions - Unresolved action**

Specifies the action that the EJB container must take if resources are uncommitted by an application in a local transaction.

This property is an IBM extension to the standard J2EE deployment descriptor. This setting is applicable only when **Resolution control** is set to Application. A local transaction context is created when a method runs in what the EJB specification refers to as an unspecified transaction context.

|           |                                      |
|-----------|--------------------------------------|
| Data type | String                               |
| Default   | Rollback                             |
| Range     | Valid values are Commit and Rollback |

Additional information about these settings follows:

#### **Commit**

At end of the local transaction context, the container instructs all unresolved local transactions to commit.

#### **Rollback**

(Default) At end of the local transaction context, the container instructs all unresolved local transactions to roll back.

### **Local Transactions - Resolution control**

Specifies how the local transaction is to be resolved before the local transaction context ends: by the application through user code or by the EJB container.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |  |
|-----------|--|
| Data type | String   |
| Range     | Valid values are Application and ContainerAtBoundary |

Additional information about these settings follows:

### Application

When this setting is used, your code must either commit or roll back the local transaction. If this does not occur, the runtime environment logs a warning and automatically commits or rolls back the connection as specified by the **Unresolved action** setting.

### ContainerAtBoundary

When this setting is used, the container takes responsibility for resolving each local transaction. This provides you with a programming model similar to global transactions in which your code simply gets a connection and performs work within it. User code does not have to handle local transactions.

- - If the **Boundary** attribute is set to `ActivitySession`, then the local transactions are enlisted as `ActivitySession` resources and directed to complete by the `ActivitySession`.
  - If the the **Boundary** attribute is set to `BeanMethod`, then the local transactions are committed at method end by the container.

Connections are never committed automatically by the resource adapter when this value is configured for the bean **Unresolved action** is not used. An application cannot call `Connection.LocalTransaction.begin()` when using this policy and receives an exception from the resource adapter if it does so.

When using a **Resolution control** of `ContainerAtBoundary`, applications must get connection handles *after* the local transaction context boundary has been started by the container. The application should close the connection before the end of the boundary, although any work performed on the connection is not committed or rolled back until the local transaction context ends. This model of connection usage is sometimes referred to as the "get-use-close" model.

This value is supported only for EJB components that use container-managed transactions. It is not supported for web components or for enterprise beans that use bean-managed transactions.

### Local Transactions - Boundary

Specifies the duration of a local transaction context. This property does not apply to stateless session beans.

This property is an IBM extension to the standard J2EE deployment descriptor.

|           |   |
|-----------|---|
| Data type | String  |
| Default   | BeanMethod  |
| Range     | Valid values are <code>BeanMethod</code> and <code>ActivitySession</code> |

Additional information about valid settings follows:

### BeanMethod

When this setting is used, the local transaction begins when the method begins and ends when the method ends.

### ActivitySession

When this setting is used, the local transaction must be resolved within the

scope of any ActivitySession in which it was started or, if no ActivitySession context is present, within the same bean method in which it was started.

### **JNDI name**

Specifies the JNDI name of the bean's home interface. This is the name under which the enterprise bean's home interface is registered and therefore, is the name that must be specified when an EJB client does a lookup of the home interface.

### **WAS Enterprise assembly settings for session beans**

Use this page to configure Enterprise Application Server functions for session beans.

**Task references:** Task references enable developers to programmatically set the current task under which a component is executing.

**Name:** The name of the task reference. This name corresponds to the string that the developer uses to set the current task.

**Task:** The task that the run-time environment will associate with execution of the current component.

**Name:** The name of the task.

Programmatically setting the logical task name causes this task name to be associated with requests from this J2EE component.

**Description:** A description of the task.

An optional field provided as a convenience for the developer.

**Internationalization type:** The Internationalization type of a session bean indicates whether the bean or the hosting EJB container will manage internationalization context on business method invocations. For session beans the Internationalization type can be configured to **Application** or **Container**, and defaults to Container. The setting applies to all business methods of a session bean.

Select whether the session bean will employ **Application** or **Container** Internationalization context management using the **Internationalization type** drop-down list.

Default

Container

Range

### Application

Under Application-managed Internationalization (AMI) session beans are responsible to manage (set) invocation context elements using the internationalization context API. The container suspends internationalization context imported from the calling process on business method invocations; to continue propagating this context on subsequent business method requests, a method must use the API to transfer caller context elements to the invocation context scoped to that method. invocation context elements not set using the API default to the respective elements of the current process when accessed using the API or when propagated on outgoing business method requests.

Specify Application-managed Internationalization for session beans having internationalization context management requirements not supported by Container-managed Internationalization.

### Container

Under Container-managed Internationalization (CMI) session beans have read-only access to the internationalization context API and rely solely on the container to manage Invocation context in accordance with the **Run as** field of the applicable Container Internationalization attribute. On a business method invocation, the container scopes the context indicated by the **Run as** field to the method implementation; when the method has completed, the container removes this context from scope. If the method is not indicated within a Container Internationalization attribute, the container scopes the caller context by default. Invocation context elements unavailable to the container default to the respective elements of the server JVM when accessed using the API or when propagated on outgoing business method requests

---

## EJB containers

An Enterprise JavaBeans (EJB) container provides a run-time environment for enterprise beans within the application server. The container handles all aspects of an enterprise bean's operation within the application server and acts as an intermediary between the user-written business logic within the bean and the rest of the application server environment.

One or more EJB modules, each containing one or more enterprise beans, can be installed in a single container.

The EJB container provides many services to the enterprise bean, including the following:

- Beginning, committing, and rolling back transactions as necessary.
- Maintaining pools of enterprise bean instances ready for incoming requests and moving these instances between the inactive pools and an active state, ensuring that threading conditions within the bean are satisfied.
- Most importantly, automatically synchronizing data in an entity bean's instance variables with corresponding data items stored in persistent storage.

By dynamically maintaining a set of active bean instances and synchronizing bean state with persistent storage when beans are moved into and out of active state, the container makes it possible for an application to manage many more bean instances than could otherwise simultaneously be held in the application server's memory. In this respect, an EJB container provides services similar to virtual memory within an operating system.

Between transactions, the state of an entity bean can be cached. The EJB container supports option A, B, and C caching.

For more information about EJB containers, see "Resources for learning."

---

## Managing EJB containers

Each application server can have a single EJB container; one is created automatically for you when the application server is created. The following steps are to be performed only as needed to improve performance after the EJB application has been deployed.

### Steps for this task

1. **(Optional)** Adjust EJB container settings.
2. **(Optional)** Adjust EJB cache settings.

### What to do next

If adjustments do not improve performance, consider adjusting access intent policies for entity beans, reassembling the module, and redeploying the module in the application.

## EJB container settings

Use this page to configure and manage a specific EJB container.

To view this administrative console page, click **Servers > Application Servers > *serverName* > EJB Container**.

### **Passivation directory**

Specifies the directory into which the container saves the persistent state of passivated stateful session beans.

Beans are passivated when the number of active bean instances becomes greater than the cache size specified in the container configuration. When a stateful bean is passivated, the container serializes the bean instance to a file in the passivation



directory and discards the instance from the bean cache. If, at a later time, a request arrives for the passivated bean instance, the container retrieves it from the passivation directory, deserializes it, returns it to the cache, and dispatches the request to it. If any step fails (for example, if the bean instance is no longer in the passivation directory), the method invocation fails.

### Inactive pool cleanup interval

Specifies the interval at which the container examines the pools of available bean instances to determine if some instances can be deleted to reduce memory usage.

|           |                |
|-----------|----------------|
| Data type | Integer        |
| Units     | Milliseconds   |
| Range     | Greater than 0 |

### Default datasource JNDI name

Specifies the JNDI name of a data source to use if no data source is specified during application deployment. This setting is not applicable for EJB 2.x-compliant CMP beans.

Servlets and enterprise beans use *data sources* to obtain these connections. When configuring a container, you can specify a default data source for the container. This data source becomes the default data source used by any entity beans installed in the container that use container-managed persistence (CMP).

The default data source for a container is secure. When specifying it, you must provide a user ID and password for accessing the data source.

Specifying a default data source is optional if each CMP entity bean in the container has a data source specified in its configuration. If a default data source is not specified and a CMP entity bean is installed in the container without specifying a data source for that bean, applications cannot use that CMP entity bean.

### Initial state

Specifies the execution state requested when the server first starts.

|           |                                      |
|-----------|--------------------------------------|
| Data type | String                               |
| Default   | Started                              |
| Range     | Valid values are Started and Stopped |

## EJB container system properties

In addition to the settings accessible from the administrative console, you can set the following system property by command-line scripting:

### **com.ibm.websphere.ejbcontainer.poolSize**

Specifies the size of the pool for the specified bean type. This property applies to stateless, message-driven and entity beans. If you do not specify a default value, the container defaults of 50 and 500 are used.

Set the pool size for a given entity bean as follows:

```
beantype
=
min
,
max[:
beantype
```

```
=  
min  
,  
max...]
```

*beantype* is the J2EE name of the bean, formed by concatenating the application name, the # character, the module name, the # character, and the name of the bean (that is, the string assigned to the <ejb-name> field in the bean's deployment descriptor). *min* and *max* are the minimum and maximum pool sizes, respectively, for that bean type. Do not specify the square brackets shown in the previous prototype; they denote optional additional bean types that you can specify after the first. Each bean-type specification is delimited by a colon (:).

Use an asterisk (\*) as the value of *beantype* to indicate that all bean types are to use those values unless overridden by an exact bean-type specification somewhere else in the string, as follows:

```
*=30,100
```

To specify that a default value be used, omit either *min* or *max* but retain the comma (,) between the two values, as follows (split for publication):

```
SMAApp#PerfModule#TunerBean=54,  
:SMAApp#SMModule#TypeBean=100,200
```

You can specify the bean types in any order within the string.

## EJB cache settings

Use this page to configure and manage the cache for a specific EJB container. To determine the cache absolute limit, multiply the number of enterprise beans active in any given transaction by the total number of concurrent transactions expected. Then, add the number of active session bean instances. You can use the Tivoli Performance Viewer to view bean performance information.

To view this administrative console page, click **Servers > Application Servers > *serverName* > EJB Container > EJB Cache Settings**.

### Cleanup interval

Specifies the interval at which the container attempts to remove unused items from the cache in order to reduce the total number of items to the value of the cache size.

The cache manager tries to maintain some unallocated entries that can be allocated quickly as needed. A background thread attempts to free some entries while maintaining some unallocated entries. If the thread runs while the application server is idle, when the application server needs to allocate new cache entries, it does not pay the performance cost of removing entries from the cache. In general, increase this parameter as the cache size increases.

|           |                |
|-----------|----------------|
| Data type | Integer        |
| Units     | Milliseconds   |
| Range     | Greater than 0 |
| Default   | 3000           |

### Cache size

Specifies the number of buckets in the active instance list within the EJB container.

A bucket can contain more than one active enterprise bean instance, but performance is maximized if each bucket in the table has a minimum number of instances assigned to it. When the number of active instances within the container exceeds the number of buckets, that is, the cache size, the container periodically attempts to reduce the number of active instances in the table by passivating some of the active instances. For the best balance of performance and memory, set this value to the maximum number of active instances expected during a typical workload.

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | Buckets in the hash table   |
| Range     | Greater than 0. The container selects the next largest prime number equal to or greater than the specified value. |
| Default   | 2053  |

## Container interoperability

*Container interoperability* describes the ability of WebSphere Application Server clients and servers at different versions to successfully negotiate differences in native Enterprise JavaBeans (EJB) Version 1.1 finder methods support and Java 2 Platform, Enterprise Edition (J2EE) Version 1.3 compliance.

At one time, there were significant interoperability problems among WebSphere Application Server, versions 4.0.x and 3.5.x distributed, and Version 4.0.x for zSeries. The introduction of interoperable versions of some class types solved these problems for distributed versions 3.5.6, 4.0.3, and 5 as well as for zSeries Version 4.0.x.

Older 4.0.x and 3.5.x client and application server versions do not support the interoperability classes, which makes them uninteroperable with versions that use the classes. The system property `com.ibm.websphere.container.portable` remedies this situation by enabling newer versions of the application server to turn off the interoperability classes. This lets a more recent application server return class types that are interoperable with an older client.

Depending on the value of `com.ibm.websphere.container.portable`, application servers at versions 5, 4.0.3 and later, and 3.5.6 and later, return different classes for the following:

- Enumerations and collections returned by EJB 1.1 finder methods
- EJBMetaData
- Handles to:
  - Entity beans
  - Session beans
  - Home interfaces

If the property is set to `false`, application servers return the old class types, to enable interoperability with versions 3.5.5 and earlier, and 4.0.2 and earlier. If the property is set to `true`, application servers return the new classes.

Instructions for setting the `com.ibm.websphere.container.portable` property are in the release notes for versions 3.5.6 and later, and 4.0.3 and later. The following tables show interoperability characteristics for various version combinations of application servers and clients as well as default property values for each combination.

## Interoperability of Version 3.5.x client with Version 5 application server

Clients at Version 3.5.5 and earlier are not interoperable with Version 5 servers when using:

- EJBMetaData
- Enumerations returned by EJB 1.x finder methods
- Handles to entity beans

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 3.5.5 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder` to `false`. With this setting in place, the Version 5 application server uses the updated handles but returns the enumerations and collections that were used in the earlier clients.

To interoperate with Version 5 application servers, you must upgrade all Version 3.5.x clients to Version 3.5.6 or later.

## Interoperability of Version 5 client with Version 3.5.x application server

| Client at Version 5, using this function      | Application server at Version 3.5.6, property true | Application server at Version 3.5.6, property false (default) | Application server at Version 3.5.5 and earlier |
|---|--|---|---|
| EJBMetaData                                   | Does not work across domains                       | Works   | Does not work                                   |
| Handle to session bean                        | Works  | Works   | Does not work                                   |
| Handle to entity bean                         | Does not work across domains                       | Does not work across domains                                  | Does not work across domains                    |
| Enumeration returned by EJB 1.x finder method | Works  | Works   | Works   |

## Interoperability of Version 4.0.x client with Version 5 application server

Ideally, all 4.0.x clients that use Version 5 application servers should be at Version 4.0.3 or later.

Version 5 application servers return the interoperability class types by default (`true`). This can cause interoperability problems for distributed clients at versions 4.0.1 or 4.0.2. In particular, problems can occur with collections and enumerations returned by EJB 1.1 finder methods.

Although it is strongly discouraged, you can set `com.ibm.websphere.container.portable` to `false` on a Version 5 application server. This causes the application server to return the old class types, providing interoperability with clients at Version 4.0.2 and earlier. This is discouraged because:

- The Version 5 application server instance would become non-J2EE 1.3 compliant with regard to handles, home interface handles, and EJBMetaData.
- EJB 1.x finder methods return collection and enumeration objects that do not originate from `ejbportable.jar`.
- Interoperability restrictions still exist with the property set to `false`.

- Version 5 client handles to entity beans and home interfaces do not work across domains for the server you set to false.

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 4.0.2 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder` to `false`. With this setting in place, the Version 5 application server uses the updated handles but returns the enumerations and collections that were used in the earlier clients.

#### Interoperability of client at Version 4.0.2 and earlier with Version 5 application server

| Client at Version 4.0.2 and earlier, using this function | Application server at Version 5, property true (default) | Application server at Version 5, property false |
|--|--|---|
| EJBMetaData  | Does not work  | Works for 4.0.2 client                          |
| Handle to session bean                                   | Does not work  | Works   |
| Handle to entity bean                                    | Does not work  | Does not work across cells                      |
| Enumeration returned by EJB 1.x finder method            | Does not work  | Works   |
| Collection returned by EJB 1.x finder method             | Does not work  | Works   |
| Handle to home interface                                 | Does not work  | Does not work across cells                      |

If you would like to use updated Handle classes in EJB 2.x-compliant beans but have one of the older clients (versions 3.5.5 and earlier, and 4.0.2 and earlier) installed, set the system property `com.ibm.websphere.container.portable.finder` to `false`. With this setting in place, the Version 5 server uses the new Handle classes but returns the older enumeration and collection classes.

#### Interoperability of client at Version 4.0.3 and later with Version 5 application server

Clients at Version 4.0.3 and later work well with Version 5 application servers. However, if you set the `com.ibm.websphere.container.portable` to `false`, client handles to entity beans and home interfaces do not work across domains for the server you set to false.

| Client at Version 4.0.3 and later, using this function | Application server at Version 5, property true (default) | Application server at Version 5, property false |
|--|--|---|
| EJBMetaData  | Works  | Works   |
| Handle to session bean                                 | Works  | Works   |
| Handle to entity bean                                  | Works  | Does not work across cells                      |
| Enumeration returned by EJB 1.x finder method          | Works  | Works   |
| Collection returned by EJB 1.x finder method           | Works  | Works   |
| Handle to home interface                               | Works  | Does not work across cells                      |

#### Interoperability of Version 5 client with Version 4.0.x application server

Clients at Version 5 work well with Version 4.0.3 application servers if you set `com.ibm.websphere.container.portable` to `true`. Client handles to entity beans and home interfaces do not work across domains for any Version 4.0.3 server with `com.ibm.websphere.container.portable` at the default value, `false`. Version 5 client handles to application servers at Version 4.0.2 and earlier also have restrictions.

| Client at Version 5, using this function      | Application server at Version 4.0.3, property true | Application server at Version 4.0.3, property false (default) | Application server at Version 4.0.2 or earlier |
|---|--|---|--|
| EJBMetaData                                   | Works  | Works   | Works for 4.0.2 server only                    |
| Handle to session bean                        | Works  | Works   | Works  |
| Handle to entity bean                         | Works  | Does not work across domains                                  | Does not work across domains                   |
| Enumeration returned by EJB 1.x finder method | Works  | Works   | Works  |
| Collection returned by EJB 1.x finder method  | Works  | Works   | Works  |
| Handle to home interface                      | Works  | Does not work across domains                                  | Does not work across domains                   |

### Interoperability of zSeries Version 4.0.x client with Version 5 application server

The only valid configuration for container interoperability with zSeries Version 4.0.x clients is the default configuration for the Version 5 application server.

### Interoperability of Version 5 client with zSeries Version 4.0.x application server

Version 5 clients should work with a zSeries Version 4.0.x application server with the correct interoperability fixes described in the zSeries documentation. The interoperability characteristics should be the same as for a Version 4.0.3 distributed application server with the property set to `true`.

| Client at Version 5, using this function      | zSeries application server at Version 4.0.x |
|---|---|
| EJBMetaData                                   | Works                                       |
| Handle to session bean                        | Works                                       |
| Handle to entity bean                         | Works                                       |
| Enumeration returned by EJB 1.x finder method | Works                                       |
| Collection returned by EJB 1.x finder method  | Works                                       |
| Handle to home interface                      | Works                                       |

### Interoperability of the handle formats in WebSphere Application Server, Version 5 and Version 5.0.1

#### 5.0.2

Applications that attempt to persist handles to enterprise beans and **EJBHome** needed to subclass `ObjectInputStream` in WebSphere Application Server, Version 5. This action was required so that the subclass `ObjectInputStream` could utilize the context class loader to resolve the classes for enterprise beans and **EJBHome** stubs.

In addition, handles created and persisted in WebSphere Application Server, Version 5 only work with objects that have an unchanged remote interface. If the remote interface is changed, the handle is no longer valid because the stub is serialized inside the handle and its serial Version UID changes if the remote interface changes.

This release introduces a new handle persistence mechanism that avoids the implementation drawbacks of the previous version. However, if handles are used for this WebSphere Application Server deployment, you should consider the following issues when applying this update, future WebSphere Application Server Fix Packs and EJB Container cumulative fixes for WebSphere Application Server, Version 5.

If a WebSphere Application Server, Version 5 persisted handle or home handle is encountered by a WebSphere Application Server, Version 5.0.1 system, it can be read and utilized. In addition, it will be converted to WebSphere Application Server, Version 5.0.1 format if it is re-persisted. The WebSphere Application Server, Version 5.0.1 format cannot be read by a WebSphere Application Server, Version 5 system unless PQ72184 is applied.

Problems arise when handles are persisted and shared across systems that are not at the WebSphere Application Server, Version 5.0.1 level or later. However, a Version 5 system can receive a handle from Version 5.0.1 remotely through a call to get a handle on an enterprise bean or a `getHomeHandle` on an **EJBHome**. The remote call will succeed, however, any attempt to persist it on the Version 5 system will have the same limitations regarding the use of `ObjectInputStream` and changes in remote interface invalidating the persisted handle.

When your application stores handles persistently and shares this persistence with multiple clients or application servers, apply WebSphere Application Server, Version 5.0.1 or PQ72184 to both the client and server systems at the same time. Failure to do so can result in the inability of these systems to read the handle data stored by upgraded systems. Also, handles stored by the WebSphere Application Server, Version 5 can force the applications of the updated system to still subclass `ObjectInputStream`. Applications using the WebSphere Application Server Enterprise, Version 5 scheduler and process choreographer, are affected by these changes. These users should update their Version 5 systems at the same time with either Version 5.0.1 or PQ72184.

If the applications store handles in the session context, or locally in a file on the same system, that is not shared by other applications, on different systems, they might be able to update their systems individually, rather than all at once. If Client Container and thin client applications do not share persisted handle data, they can be updated as needed as well. However, handles created and persisted in WebSphere Application Server, Version 5, Version 4.0.3 and later (with the property flag set), or Version 3.5.7 and later (with the property flag set) are not usable if either the home or the remote interface changes.

If any WebSphere Application Server, Version 3.5.7 or Version 4.0.3 and later enables the system property `com.ibm.websphere.container.portable` to **true**, any handles to objects on that server have the same interoperability limitations. In

addition, if any WebSphere Application Server, Version 3.5.7 and later or Version 4.0.3 applications store a handle obtained from a WebSphere Application Server, Version 5 or Version 5.0.1, the same restrictions apply, regarding the need to subclass `ObjectInputStream` and the usability of handles after a change to the remote interface is made.

### **Replication of the Http Session and Handles**

This note applies to you if you place Handles to Homes or EJBs, or EJB or EJBHome references in the Http Session in your application and you use Http Session Replication. If you intend to replicate a mixed environment of Version 5.0.0 and Version 5.0.1 or 5.0.2 machines you should first apply the latest Version 5.0.0 container cumulative e-fix to the Version 5.0.0 machines before allowing the Version 5.0.1 or 5.0.2 server into the typology. The reason for this is that Version 5.0.0 servers are not able to understand the persisted Handle format used on the Version 5.0.1 and 5.0.2 server. This is similar to the case of Version 5.0.0 and Version 5.0.1 or 5.0.2 systems trying to use a shared database, mentioned above. But in this case, it is the Http Session object and not the database providing the persistence.

### **Top Down Deployment Mapping**

The size of the Handle objects has grown due to the fix put in to allow serialization and deserialization to occur without the previous requirements of subclassing the `ObjectInputStream` and so on. Top down deployment of an object that contains EJB and EJBHome references create a database table ddl that has a field of 1000 bytes of VARCHAR for BITDATA which will contain the Handle. It might be that your object's Handle does not fit in the 1000 byte default field, and you might need to adjust this to a higher value. You might try increments of 250 bytes, that is, 1250, 1500, and so on.

---

## **Deploying EJB modules**

### Before you begin

Assemble one or more EJB modules, (assemble one or more Web modules), and (assemble them into a J2EE application).

### Steps for this task

1. Prepare the deployment environment.
2. (Deploy the application.)
3. Update the configuration for each EJB module as needed for the deployment environment.
4. For information about the EJB deployment tool, see the [ ].

### What to do next

The next step is to test and debug the module.

## **EJB module collection**

Use this page to manage the EJB modules deployed in a specific application.



To view this administrative console page, click **Applications > Enterprise Applications > *applicationName* > EJB modules**. Click the check boxes to select one or more of the EJB modules in your collection.

### URI

When resolved relative to the application URL, this specifies the location of the module's archive contents on a file system. The URI matches the <ejb> or <web> tag in the <module> tag of the application deployment descriptor.

## EJB module settings

Use this page to configure and manage a specific deployed EJB module.

To view this administrative console page, click **Applications > Enterprise Applications > *applicationName* > EJB modules > *moduleName***.

### URI

When resolved relative to the application URL, this specifies the location of the module archive contents on a file system. The URI must match the URI of a ModuleRef URI in the deployment descriptor of the deployed application (EAR).

### Alternate DD

Specifies a deployment descriptor to be used at run time instead of the one installed in the module.

### Starting weight

Specifies the order in which modules are started when the server starts. The module with the lowest starting weight is started first.

|           |                |
|-----------|----------------|
| Data type | Integer        |
| Default   | 5000           |
| Range     | Greater than 0 |

---

## Enterprise beans: Resources for learning

Use the following links to find relevant supplemental information about enterprise beans. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to this product but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Planning, business scenarios, and IT architecture
- Programming model and decisions
- Programming instructions and examples
- Programming specifications

### Planning, business scenarios, and IT architecture

•



### Mastering Enterprise JavaBeans

(<http://www.theserverside.com/books/masteringEJB/index.jsp>)

A comprehensive treatment of Enterprise JavaBeans (EJB) programming in nonprintable form (PDF). One must be registered to download the PDF, but registration is free. Information about purchasing a hardcopy is available on the Web site.

- *Enterprise JavaBeans* by Richard Monson-Haefel (O'Reilly and Associates, Inc.: Third Edition, 2001)

### Programming model and decisions

- 



**Read all about EJB 2.0** (<http://www-106.ibm.com/developerworks/java/library/j-jw-ejb20/>)

A comprehensive overview of the specification.

- 



**The J2EE Tutorial** ([http://java.sun.com/j2ee/tutorial/1\\_3-fcs/index.html](http://java.sun.com/j2ee/tutorial/1_3-fcs/index.html))

This set of articles by Sun Microsystems covers several EJB-related topics, including the basic programming models, persistence, and EJB Query Language.

### Programming instructions and examples

- 



#### **Rules and Patterns for Session Facades**

([http://www7b.boulder.ibm.com/wsdd/library/techarticles/0106\\_brown/sessionfacades.html](http://www7b.boulder.ibm.com/wsdd/library/techarticles/0106_brown/sessionfacades.html))

EJB programming practice: Fronting entity beans with a session-bean facade.

- 



#### **WebSphere Application Server Development Best Practices for Performance and Scalability**

([http://www-4.ibm.com/software/webservers/appserv/ws\\_bestpractices.pdf](http://www-4.ibm.com/software/webservers/appserv/ws_bestpractices.pdf))

Programming practice for enterprise beans and other types of J2EE components.

- 



#### **Optimistic Locking in IBM WebSphere Application Server 4.0.2**

(<http://www7b.boulder.ibm.com/wsdd/>)

Examples of the effect of optimistic concurrency on application behavior. Although the paper is based on a previous version of this product, the data access issues discussed in it are current.

This paper does not seem to be available directly by URL. To view this paper, visit the specified URL and search on "optimistic locking"

## Programming specifications

- 



**What's new in the Enterprise JavaBeans 2.0 Specification?**  
(<http://java.sun.com/products/ejb/2.0.html>)

You can also download the specification itself from this URL.

- 



**Java™ 2 Platform: Compatibility with Previous Releases**  
(<http://java.sun.com/products/jdk/1.3/compatibility.html>)

This Sun Microsystems article includes both source and binary compatibility issues.

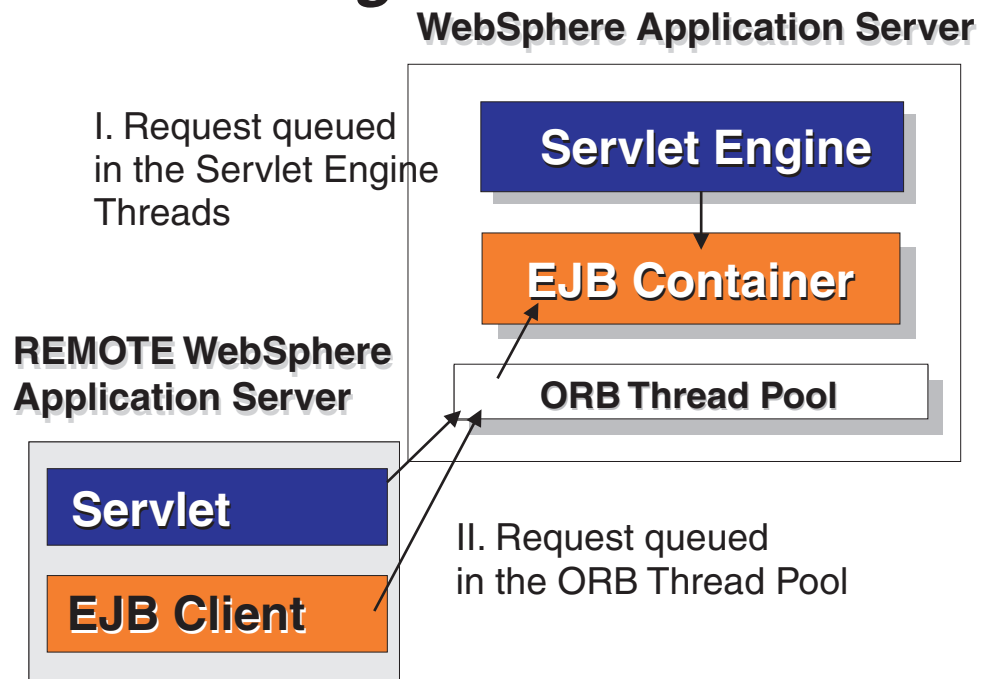
---

## EJB method Invocation Queuing

Method invocations to enterprise beans are only queued for remote clients, making the method call. An example of a remote client is an enterprise Java bean (EJB) client running in a separate Java virtual machine (JVM) (another address space) from the enterprise bean. In contrast, no queuing occurs if the EJB client, either a servlet or another enterprise bean, is installed in the same JVM on which the EJB method runs and on the same thread of execution as the EJB client.

Remote enterprise beans communicate by using the Remote Method Invocation over an Internet Inter-Orb Protocol (RMI-IIOP). Method invocations initiated over RMI-IIOP are processed by a server-side object request broker (ORB). The thread pool acts as a queue for incoming requests. However, if a remote method request is issued and there are no more available threads in the thread pool, a new thread is created. After the method request completes the thread is destroyed. Therefore, when the ORB is used to process remote method requests, the EJB container is an open queue, due to the use of unbounded threads. The following illustration depicts the two queuing options of enterprise beans.

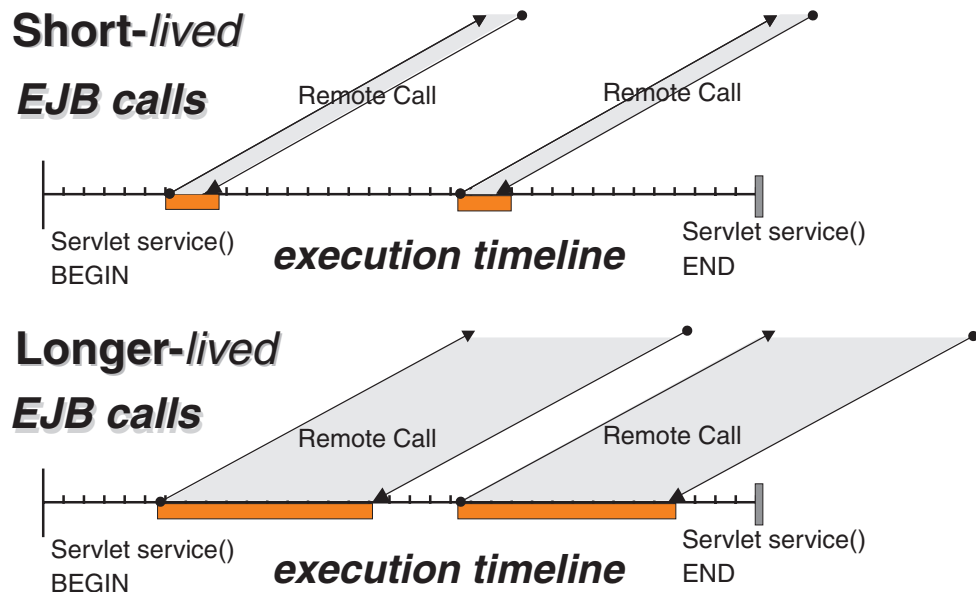
# EJB Queuing



The following are two tips for queuing enterprise beans:

- Analyze the calling patterns of the EJB client.

When configuring the thread pool, it is important to understand the calling patterns of the EJB client. If a servlet is making a small number of calls to remote enterprise beans and each method call is relatively quick, consider setting the number of threads in the ORB thread pool to a value lower than the Web container thread pool size value.



The degree to which the ORB thread pool value needs increasing is a function of the number of simultaneous servlets, that is, clients, calling enterprise beans and the duration of each method call. If the method calls are longer or the

applications spend a lot of time in the ORB, consider making the ORB thread pool size equal to the Web container size. If the servlet makes only short-lived or quick calls to the ORB, servlets can potentially reuse the same ORB thread. In this case, the ORB thread pool can be small, perhaps even one-half of the thread pool size setting of the Web container.

- **Monitor the percentage of configured threads in use.**

Tivoli Performance Viewer shows a metric called *percent maxed*, which is used to determine how often the configured threads are used. A value that is consistently in the double-digits, indicates a possible bottleneck at the ORB. Increase the number of threads.



---

## Chapter 5. Using extended messaging in applications

These topics provide information about implementing WebSphere enterprise applications that use extended messaging.

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface. *Extended messaging* extends the base JMS support, support for EJB 2.0 message-driven beans, and the Enterprise Java Bean (EJB) component model, to use the existing container-managed persistence and transactional behavior.

Using extended messaging, you can build enterprise beans that can provide messaging services along with methods that implement business logic. The enterprise beans can use the standard JMS styles of messaging (point-to-point and publish/subscribe). However, with extended messaging, the JMS usage is simplified, because JMS support is managed by the extended messaging service. This helps to effectively separate business logic from the messaging infrastructure. The use of data mapping enables messages to drive existing or new enterprise beans as though they were invoked from any EJB client.

You can use WebSphere Studio Application Developer Integration Edition to develop applications that use extended messaging. You can use the WebSphere Application Server runtime tools, like the administrative console, to deploy and administer applications that use extended messaging.

For more information about implementing WebSphere enterprise applications that use extended messaging, see the following topics:

- "Extended messaging - overview"
- "Extended messaging - application usage scenarios"
- "Extended messaging - components"
- "Designing an enterprise application to use extended messaging"
- "Developing an enterprise application to use extended messaging"
- "Deploying an enterprise application to use extended messaging"
- "Configuring extended messaging service resources"
- "Troubleshooting extended messaging"

---

### Extended messaging - overview

*Extended messaging* extends the base JMS support, support for EJB 2.0 message-driven beans, and the Enterprise Java Bean (EJB) component model, to use the existing container-managed persistence and transactional behavior.

In addition to providing such *container-managed messaging*, extended messaging provides new types of enterprise beans and administrative objects for messaging, and new functionality like data mapping and late response handling. (The abbreviation, CMM, for the term *container-managed messaging* is sometimes used to represent extended messaging.)

Extended messaging uses the bean-managed messaging implementation to provide the JMS interfaces, which ensures that both bean-managed and extended messaging use consistent JMS support.

An application that uses extended messaging can receive messages by using a *receiver bean*, either by the `onMessage()` method of a message-driven bean or by a stateless session bean that polls for a message from a named destination. With extended messaging and a message-driven bean, code within the bean can use the message to invoke business logic, as either a method within the same bean or another enterprise bean. Both the incoming message and the invocation of the receiver bean can be included within the scope of a transaction. For outbound messages, an application calls a *sender bean* that turns a method call into a JMS message that is then sent asynchronously. These message beans are implemented as enterprise beans by WebSphere Application Server. Application developers can create these message beans by using WebSphere Studio Application Developer Integration Edition, although they can be created without the help of WebSphere Studio.

With extended messaging, the JMS usage is simplified, because JMS support is managed by the extended messaging service. This helps to effectively separate business logic from the messaging infrastructure. Also, the use of data mapping enables messages to drive existing or new enterprise beans as though they are invoked from any EJB client. WebSphere Studio enables the types of message beans that support extended messaging to be developed easily and hides the messaging infrastructure from developers.

For more conceptual information about extended messaging, see the following topics:

- "Extended messaging - application usage scenarios"
- "Extended messaging - components"
- "Extended messaging - receiving messages"
- "Extended messaging - sending messages"
- "Extended messaging - data mapping"
- "Extended messaging - handling late responses"
- "Extended messaging - transactional support"
- "Extended messaging - exception handling"


## Extended messaging - receiving messages

To receive messages, applications that use extended messaging use a receiver bean, which can be a message-driven bean or a session bean:

- A receiver bean (deployed as a message-driven bean) is invoked when a message arrives at a JMS destination for which a listener is active.
- An application-callable receiver bean (deployed as a session bean) polls a JMS destination until a message arrives, gets the parsed message as an object, and can use getter methods to retrieve the message data.

**Receiving messages with extended messaging.** This figure shows an application calling a receiver bean (as a session bean) to receive messages from the JMS destination defined on an input port. The application also calls the `ReplySender()` method of the receiver bean to send a reply to the original message received. For more information about what is shown in this figure, see the text that accompanies this figure.





When a receiver bean gets a message, it can invoke another method passing either the JMS Message, or a set of parameters extracted from the message content. The invoked method can be contained in the receiver bean or in another enterprise bean (which is the preferred application structure). If data mapping is used, the method invoked by a receiver bean is unaware of the original JMS message.

In addition to receiving messages, extended messaging enables applications to reply to received messages in either of the following ways:

- **Sending a synchronous reply.** In this mode, which can be used by only receiver beans deployed as message-driven beans, the reply from the method invoked by the receiver bean is mapped to a JMS message and sent as a reply to the original message, using the replyTo field in the JMS header as the target destination.
- **Sending an asynchronous reply.** In this mode, which cannot be used receiver beans deployed as message-driven beans, the application calls the ReplySender() method to send the reply message. If the reply is passed as a set of parameters to the ReplySender() method, the reply is mapped to a JMS message before being sent.

If a receiver bean gets a JMS message, then depending upon the programming model (associated with the receiver bean by WebSphere Studio), one of the following interactions occurs:

- **Receive a request and send no reply.**  
The receiver bean provides an anonymous invocation in the form of a method call. The data passed to the method is either the JMS message (if no data mapping is invoked) or a set of parameters mapped from the JMS message. The receiver bean cannot return a reply to this invocation. This mode of interaction can be used with point-to-point or publish/subscribe messaging.
- **Receive a request and send a synchronous reply.**  
If the receiver bean gets a message, it invokes another method either in the same bean or another enterprise bean. When the method returns, the data returned from that call is mapped to a JMS message and sent to the reply destination specified in the original request message. The type of reply destination (queue or topic) must be the same as the type used by the original request received.
- **Receive a request and send an asynchronous reply.**  
For a receiver bean deployed as a message-driven bean, the reply is returned (using the bean's ReplySender method) to the replyTo destination specified in the original request message.


For a receiver bean deployed as a session bean, the reply is returned to the destination defined in the input port for the receiver bean.

In addition to the asynchronous model of this interaction, this mode of interaction enables a method to send multiple replies to a single invocation.

## Extended messaging - sending messages

To send messages, applications that use extended messaging call a method on a *sender bean*. A sender bean turns its method invocation into a JMS message, then passes that message to JMS. If needed, the sender bean can retrieve a response message, then translate that message into a result value and return it to the caller. If data mapping is used, the method that invokes a sender bean is unaware of the original JMS message. The sender bean methods can use data mapping to build JMS messages from data passed on the method call.

**Sending messages with extended messaging.** This figure shows an application calling a sender bean to send messages to the JMS destination defined on an output port. The application also calls the `receiveResponse()` method of the sender bean to receive a reply to the original message sent. For more information about what is shown in this figure, see the text that accompanies this figure.



A sender bean is an enterprise bean (stateless session bean) that can be built by WebSphere Studio Application Developer. A sender bean should not contain any application logic, to help preserve the separation between the messaging and business logic.

Each method defined on a class that implements a sender bean has one of the following modes of interaction (which is defined when the sender bean is built). The interaction extends the sender interface to address the issue of synchronizing anonymous invocations.

- **Send a request and receive no response.**

To send a JMS message, an application invokes the sender bean's method. The caller of the sender bean's method cannot receive a response to the message sent. This mode of interaction can be used with point-to-point or publish/subscribe messaging.

- **Send a request and receive a synchronous response.**

To send a JMS message and wait for a synchronous response, an application invokes the sender bean's method. The sender bean uses the message sender (an interface to JMS provided by extended messaging) to send the message and,

when the response is received, to return the response message to the caller of the sender bean. This mode of interaction can be used with point-to-point messaging only.

- **Send a request and receive a deferred response.**

To send a JMS message and wait for a deferred response, an application invokes the sender bean's method. The sender bean uses the message sender to send the message, then returns to the caller without waiting for the response. The response is returned by a generated `receiveResponse()` method. This mode of interaction enables an application to receive more than one response message, as the application is responsible for retrieving the responses. This mode of interaction can be used with point-to-point messaging only.

## Extended messaging - data mapping

A message bean can use data mapping to map between a JMS message and data as arguments:

- With data mapping, the target method of a receiver bean for an anonymous invocation receives the contents of an asynchronous message as arguments. The extended messaging service parses the JMS message and maps from the JMS message to the method arguments. Similarly, to send a message, an application invokes a method on a sender bean with appropriate arguments. The extended messaging service packs appropriate arguments into a JMS message then sends the asynchronous request.
- Without data mapping, the target method of a receiver bean for an anonymous invocation receives a JMS message; no data mapping is performed by extended messaging. Similarly, to send a message, an application invokes a method of a sender bean with a JMS message.

If a developer selects data mapping when creating a sender or receiver bean, extended messaging uses the parameter properties specified on the sender or receiver bean method signatures to perform the data mapping.

## Extended messaging - handling late responses

If an application uses a sender bean to send a message, it can optionally retrieve a response to the message. The sender bean can either wait for the response or defer retrieval of the response. Sometimes a response is delayed within the messaging infrastructure, and therefore the application cannot receive the response. Extended messaging can retrieve such a response message (referred to as a *late-response* message) when it does arrive and pass it to a message-driven bean provided by the application to handle late responses. The message-driven bean used to handle the late response is a standard EJB 2.0 message-driven bean or a receiver bean deployed as a message-driven bean. The deployed message-driven bean can then perform its processing on the message.

Late responses should not be considered normal application behaviour.

For extended messaging to handle late responses for an application, the sender bean must be deployed with the **Handle late responses** option enabled.

### Definition of a late response

A late response occurs when the application is no longer able to retrieve responses to messages that it has sent, as follows:

- **Send with deferred response.**

The application (enterprise bean) repeatedly tries to retrieve a response until it ends. When the application no longer wants to retry to get a response, it can register a request for extended messaging to handle the late response, by calling a `registerLateResponse()` method on the sender bean.

- **Send with synchronous response handling.**

When the sender bean sends a message, it waits for the response. The result of this is that either the sender bean retrieves the response message or a timeout error occurs. If the system raises a timeout error, the application can no longer retrieve a response to the message. At this time the extended messaging service registers the the message for a late response.

## Handling responses

Extended messaging handles responses in the following stages:

1. Registering an interest in having a late response retrieved when it is available.

To request the system to handle late responses for a sender bean, you deploy the sender bean with the **Handle late responses** extension to the deployment Descriptor.

If selected, the **Handle late responses** option defines that extended messaging should pass the response, when it becomes available, to the message-driven bean provided by the application to handle late responses. When the sender bean is deployed a specialized listener port is associated with the bean. This listener port is known as a *handle late response listener port*.

If the option is not selected, then the system does not handle late responses, and it is the application's responsibility to handle any late responses.

2. Starting a JMS listener to retrieve the message when it is available, which then drives the message bean to handle the JMS message.

The listener port must be defined with the following properties:

- The same JMS destination as specified as the JMS response destination on the output port used by the sender bean.
- A listener port extension with Handle Late Responses enabled.

You cannot use a temporary destination for late responses.

3. If a request is made to handle a late response, the extended messaging service immediately registers a `LateResponse` message request with the extended message consumer for the given listener port. The message request is registered independently of any transaction context that the sender bean has. A request record (containing the `MessageID` of the late response) is added to the `AsyncMessageLog` log. When the message is eventually received, it is passed to the message-driven bean deployed against the specified late response `ListenerPort`.

## Extended messaging - transactional support

The global transaction context is not flowed on asynchronous (messaging) requests, so the receipt of an asynchronous message cannot be part of some existing remotely-established transaction. Reliability in an asynchronous environment is built on the message provider's ability to guarantee a once-and-once-only message delivery.

Transactional support with extended messaging builds on, and extends, the transactional support provided with bean-managed messaging, as follows:

- Transactional support for receiving messages (receiver beans)

- Transactional support for sending messages (sender beans)

### **Transactional support for receiving messages (receiver beans)**

The extended messaging transactional behaviour for receiver beans depends upon whether the bean is a receiver bean or an application-callable receiver bean.

- For a receiver bean (deployed as a message-driven bean), incoming message receipts are defined by the Transaction attribute of the `onMessage()` bean method. Message-driven beans can use bean-managed transactions (BMT) or container-managed transactions (CMT). For message-driven beans using CMT there are only two supported transaction options: Not supported and Required. If a message is to be received within a transaction, the message-driven bean must be using CMT with the Transaction attribute set to Required.
- For a receiver bean as a session bean, the bean only supports container-managed transactions, and the behaviour is defined by the Transaction attribute of the receiver bean method.

**Dealing with retries:** In the asynchronous environment of transaction processing, rolling back a message receipt means that the message is not removed from the source destination. Although this behaviour is desirable and correct, it causes the message retained on the source queue to be reprocessed until the transaction commits. For receiver beans, you can control this behaviour as follows:

- **Receiver bean.**

To limit the number of times that a transaction is retried, you can either rely on the facilities of the JMS provider or use the retry limit facility of the Message Listener:

- WebSphere MQ JMS support provides the ability to move the message to a backout queue and uses two queue attributes, the backout threshold and the backout-requeue queue, to perform this.
- The Message Listener retry count can be used to stop the listener processing the queue if the threshold is reached. The listener behaviour can be disabled by setting the retry count value higher than the JMS provider threshold value.

- **Application-callable receiver bean.**

To limit the number of times that a transaction is retried depends on the facilities of the JMS provider to move the message to a backout queue.

### **Transactional support for sending messages (sender beans)**

The transactional behaviour for sending messages is defined by the Transaction attribute on the `send` method within the sender bean.

If the `send()` method is part of a transaction, then the sending of an outgoing message occurs within any currently active transaction. This means that the message is not transmitted until the transaction is committed. If no transaction is active when the request to send the message occurs, then the message is transmitted immediately.

The transactional behaviour where the mode of interaction for a sender bean method specifies a response (that is, either *Send message and receive synchronous response* or *Send message and receive deferred response*) depends upon the type of response, as follows:

#### **Transactional behaviour for a synchronous response**

The sending of the request message and the receipt of the response

message cannot be performed inside a transaction, because they are both performed within the same method. Therefore, the send is always non-transactional, regardless of the transactional setting for the method. The receive is either transactional or not, depending upon the Transaction setting of the method.

#### **Transactional behaviour for a deferred response**

The response message is received by a `receiveResponse()` method, which can have a different transactional behaviour to the sender method inside the sender bean. The transaction containing the send command must commit before the response can be received. The transactional behaviour is specified on the send and receive methods of the Sender bean.

## **Extended messaging - exception handling**

Extended messaging provides the following exception handling for receiver beans and sender beans:

- Error handling for receiver beans
- Error handling for sender beans

#### **Error handling for receiver beans**

The following error conditions can lead to extended messaging exceptions

- Formatting error parsing the message, when performing data mapping
- Exception thrown by the application method
- `CMMException` when sending the reply

Errors are always logged. If the application does not catch the exception, the default behavior is to roll back any active transaction. If the received message is rolled back, then it can be processed again. This can occur a number of times until the message causing the error is removed from the queue by the JMS provider. (For more information, see *Dealing with retries* in Transactional support with extended messaging.

With data mapping, if a receiver bean is deployed as a message-driven bean and a `replyTo` destination is configured, then error messages are sent as replies to that destination.

Application enterprise beans that call receiver beans deployed as session beans need to handle the `CMMException` exception. `CMMException` is an application exception which is declared in the `throws` clause of the methods in the generated receiver bean.

#### **Error handling for sender beans**

The following error conditions can lead to extended messaging exceptions

- Constructing the JMS message when data mapping from the parameters to the message
- Creating a message sender and sending the message
- Getting the response and parsing the message content

Errors are always logged. If the application does not catch the exception, the default behavior is to roll back any active transaction.

Application enterprise beans that call sender beans need to handle the `CMMEException` exception, which is declared in the `throws` clause of the methods in the generated sender bean.

---

## Extended messaging - application usage scenarios

Applications can use extended messaging to receive and send messages in a variety of ways:

- To receive messages, applications that use extended messaging use a *receiver bean* (deployed as a message-driven bean) or an *application-callable receiver bean* (deployed as a session bean):
  - A receiver bean is invoked when a message arrives at a destination for which a listener is active.
  - An application-callable receiver bean polls a destination (defined by an *input port*) until a message arrives or a timeout occurs.

In addition to receiving messages, extended messaging enables applications to send replies in response to the received messages.

For more conceptual information about receiving messages, see "Extended messaging - receiving messages".

- To send messages, applications that use extended messaging call *sender bean* methods. The sender bean sends messages to the target destination defined by an *output port*. The sender bean methods can be passed either a JMS message or a number of parameters that are mapped by extended messaging into a JMS message. Whether or not data mapping is used is specified when the application is developed.

In addition to sending messages, applications can choose to receive a response to the message, and can handle any responses either synchronously or asynchronously. If a response is not received in time, then the system can handle the late response by directing the message to a message-driven bean.

For more conceptual information about sending messages, see "Extended messaging - sending messages".

- Applications can combine receiving and sending messages in a variety of different ways. For example, a receiver bean deployed as a message-driven bean can forward the message by calling a sender bean. The receiver bean can give message data to sender bean in either of the following ways:
  - The receiver bean can pass the JMS message to the sender bean, which forwards that message.
  - The receiver bean can extract data from the initial message and send that data to the sender bean. The sender bean can then map the data values to a new JMS message, which it forwards.

The application can receive a response to the message that it sent, and then can send the message received or a new message built from data in the message received, as a response to the original message.

Also, data mapping can be used to hide the JMS message structure from the application logic. For more information about data mapping, see "Extended messaging - data mapping"

---

## Extended messaging - components

Extended messaging builds on the base support for JMS messaging and message-driven beans provided by WebSphere Application Server. The new messaging components for extended messaging are referred to as the *Message Bean package*.

### Components for receiving messages

The following components, shown in the figure Components for receiving messages, are used to receive messages:

#### Receiver bean

An application that uses extended messaging can receive messages by using a *receiver bean* (using the `onMessage()` method of a message-driven bean) or an *application-callable receiver bean* (a stateless session bean that polls for a message from a named destination). Both receiver beans and application-callable receiver beans can receive and process asynchronous messages, and optionally return selected data as a response message.

#### Input port

An input port specifies the properties needed by receiver beans as session beans, by defining the following information:


- Information about the source destination for the message to be received
- Information about how to select and handle the message received
- Optional information about a reply destination, which is used if a reply is expected and `replyTo` information is not present in the `JMSHeader` of the message received.

A receiver bean as a deployed message-driven bean uses the associated listener port, so does not need an input port. For more information about message-driven beans and listener ports, see "Message-driven beans - components" (not in this document).

For more conceptual information about receiving messages, see "Extended messaging - receiving messages".

**Components for receiving messages.** This figure shows an application calling a receiver bean (as a session bean) to receive messages from the JMS destination defined on an input port. The application also calls the `ReplySender()` method of the receiver bean to send a reply to the original message received. For more information about what is shown in this figure, see the text that accompanies this figure.





### Components for sending messages

The following components, shown in the figure Components for sending messages, are used to send messages:

#### Sender bean

Bean instances of a sender bean (also known as a *message sender bean*) can send asynchronous messages. The sender bean methods can be passed either a JMS message or a number of parameters that are mapped by extended messaging into a JMS message, which is then passed to JMS.


#### Output port

An output port specifies the properties needed by sender beans, to define the destination for the message being sent, and other optional properties if a response is expected. It is associated with the Sender Bean at deployment time and contains the following information:

- Information about the target destination for the message to be sent
- Information about how to select and handle the message to be sent
- Information about the destination used for the response.

For more conceptual information about sending messages, see Sending messages with extended messaging.

**Components for sending messages.** This figure shows an application calling a sender bean to send messages to the JMS destination defined on an output port. The application also calls the `receiveResponse()` method of the sender bean to receive a reply to the original message sent. For more information about what is shown in this figure, see the text that accompanies this figure.



## Designing an enterprise application to use extended messaging

This topic describes things to consider when designing an enterprise application to use extended messaging.

The design of JMS-usage for applications that use extended messaging is the same as the design for JMS and message-driven beans, except that the JMS-usage is simplified because JMS support is managed by the extended messaging service. For design considerations related to JMS and message-driven beans, see the following topics:

- Designing an enterprise application to use JMS
- Designing an enterprise application to use message-driven beans

The extra design consideration for applications that use extended messaging are as follows. For more detail, see the related topics.

### Steps for this task

1. For a receiver bean, decide whether to use a message-driven bean or stateless session bean.

#### Message-driven bean

You can use a deployed message-driven bean as a receiver bean, to automatically handle messages received at the associated listener port. As with any message-driven bean, when a message is received on the JMS destination monitored by the listener port, the message is passed to the `onMessage()` method of the message-driven bean.

You need to develop and deploy the message-driven bean, and configure its associated listener port, separately from the extended messaging tasks.

#### Stateless session bean

You can use a stateless session bean as a receiver bean, to poll for messages on a named destination associated with an input port.

You need to develop and deploy the session bean separately from the extended messaging tasks, but configure the associated input port as part of the extended messaging tasks.

2. Decide whether or not you want to use data mapping.

If you call the methods of sender and receiver beans with data arguments, you need to use data mapping to construct the JMS messages needed. For data mapping, you need to decide what data arguments need to be specified as properties on the sender or receiver bean method signatures.

For a receiver bean deployed as a message-driven bean, you can define the mapping behavior if a data exception is caught by extended messaging. That is, you define whether a message should be flowed back if a ReplyTo destination is defined in the JMS message header.

3. Decide whether or not you want to handle late responses.

A sender bean can optionally retrieve a response to messages sent. If a response is delayed within the messaging infrastructure, the bean cannot receive the response. Extended messaging can retrieve such a response message (referred to as a late-response message) when it does arrive and pass it to a message-driven bean provided by the application to handle late responses. To handle late responses, you need to develop and deploy a standard EJB 2.0 message-driven bean that contains a registerLateResponse() method, and associate it with a listener port to be used to receive late responses.

---

## Developing an enterprise application to use extended messaging

This topic describes how to develop an enterprise application to use extended messaging.

This task description assumes that developers are using the WebSphere Studio Application Developer to develop the application code (receiver and sender beans).

To develop an enterprise application to use extended messaging, complete the following steps:

### Steps for this task

1. Creating the Enterprise Application project.

Because the sender and receiver beans used for extended messaging are EJB 2.0 enterprise beans, you must first have created a J2EE 1.3 Enterprise Application project for which extended messaging beans will be created.

- a. Ensure that you have selected 1.3 as the highest J2EE version that is to be used in WebSphere Studio.

For example: **Window-> Preferences... J2EE preferences-> Select the highest J2EE version that is to be used-> 1.3**

- b. Create a J2EE 1.3 Enterprise Application project, as described in the WebSphere Studio article "Creating an Enterprise Application project".

2. Creating the application code.

To create the application code, use WebSphere Studio to generate the sender and receiver beans needed by the application, by completing one or more of the following subtasks as described in the WebSphere Studio Extended Messaging documentation:

- "Creating a sender bean"
- "Creating a receiver bean"
- "Creating an application-callable receiver bean"
- "Creating a sender bean and receiver bean"
- "Creating a sender bean and application-callable receiver bean"

The result of this stage is an enterprise bean, containing code automatically generated for extended messaging, that can be assembled into an .EAR file for deployment.

3. Assembling and packaging the application for deployment.

You can use WebSphere Studio to assemble and package the application for deployment.

The following aspects are specific to extended messaging:

- a. **(Optional)** Configure a message selector for a receiver bean.
- b. Associate the JNDI names for sender and receiver beans with output and input ports.
- c. **(Optional)** Specify the timeout for a sender bean response.
- d. **(Optional)** Configure that a sender bean is to handle late responses and identify the listener port to be used for late responses.

### Results

The result of this task is an .EAR file, containing an application enterprise bean with code for extended messaging, that can be deployed in WebSphere Application Server.

### What to do next

For information about deploying an application to use extended messaging, see "Deploying an enterprise application to use extended messaging"

---

## Deploying an enterprise application to use extended messaging

This topic describes how to deploy an enterprise application to use extended messaging.

This task description assumes that you have an .EAR file, which contains an application enterprise bean with code for extended messaging, that can be deployed in WebSphere Application Server.

The Application Install task is also a standard WebSphere Application Server task. As part of the install procedure you need to associate the Input and Output ports defined in System Management with the installed .EAR.

To deploy an enterprise application to use extended messaging, complete the following steps:

### Steps for this task

1. Use the administrative console to define and configure the extended messaging resources to be used by the application, as described in "Configuring extended messaging service resources".

You should define the input ports for receiver beans, the output ports for sender beans, and listener port extensions for any sender beans that are to handle late responses.

2. Ensure that the deployment descriptor attributes for the sender and receiver beans match those of the extended messaging resources that you configured using the administrative console.

The deployment descriptor values can be set when you generate the deployment code for the application. You can change the deployment descriptor values by using the application assembly tool, as described in the following topics:

- "Configuring deployment attributes for a receiver bean"
  - "Configuring deployment attributes for a sender bean".
3. If a sender bean is to handle late responses, deploy the message-driven bean to be used for late responses.

For more information about deploying message-driven beans, see "Deploying an enterprise application to use message-driven beans" (not in this document).

4. Install the application into WebSphere Application Server.

This stage is a standard WebSphere Application Server task, as described in "Installing a new application" (not in this document).

When you install the application, you are prompted to specify the name of the listener port that the application is to use for late responses. Select the listener port, then click **OK**.

## Configuring deployment attributes for a receiver bean

Use this task to configure the deployment attributes for a receiver bean for use with the extended messaging service.

You can specify these deployment attributes on each EJB method, as part of the deployment of the receiver bean. Changes to the deployment attributes override the values defined when the receiver bean was developed and deployment code was generated for the application.

**Note:** After deployment code has been generated for an application, the deployable archive is renamed with the prefix `Deployed_`. Any subsequent changes to the archive from within the Application Assembly Tool are applied to the version of the archive that existed prior to code generation. To see changes reflected in your application, you must regenerate deployment code and re-install the deployable archive.

To configure the deployment attributes for a receiver bean, you can use the Application Assembly Tool to complete the following steps:

### Steps for this task

1. Launch the Application Assembly Tool.
2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, click **File-> Open** then select the EAR file.
3. In the navigation pane, select the receiver bean instance; for example, expand *ejb\_module\_instance-> Extended messaging* then select the bean instance.  
A property dialog notebook for the receiver bean is displayed in the property pane.
4. In the property pane, specify appropriate deployment attributes:

#### **Input port**

For an application-callable receiver bean, this is the name of the input port to be used to receive messages.

### Message selector

For an application-callable receiver bean, this is a string used to select messages to be received.

For more information about the deployment attributes for receiver beans, see Extended messaging assembly properties for enterprise beans.

5. To apply the changes and close the Application Assembly Tool, click **OK**. Otherwise, to apply the values but keep the property dialog open for additional edits, click **Apply**.
6. **(Optional)** To see changes reflected in your application, regenerate deployment code and reinstall the deployable archive.

### Extended messaging assembly properties for EJB modules

Use this page to configure extended messaging settings for methods of an enterprise bean.

**Name:** Specifies a name for the mapping between extended messaging settings and one or more methods.

Datatype                      String

**Description:** Contains text that describes the mapping

Datatype                      String

**Methods:** The methods to which these settings apply.

To add a new method, click **Add**. Expand the tree to select the method or methods from the EJB module

**Reply timeout:** The time in milliseconds after which replies, delayed within the messaging infrastructure, are considered as late.

Type the global reply timeout to be used if a reply timeout is not specified on a sender method call. If you leave this field blank, no global timeout is defined.

This is the time in milliseconds after which replies, delayed within the messaging infrastructure, are considered as late and therefore the application cannot receive the response. The extended messaging service can retrieve such a response message (referred to as a *late-reply* message) when it arrives and pass it to a message bean provided by the application.

|           |  |
|-----------|--|
| Data type | Integer  |
| Units     | Milliseconds   |
| Default   | 0  |
| Range     | An integer number of milliseconds, greater than or equal to 0 (0 indicates that reply messages never timeout). |

**Message selector:** The JMS message selector to be used to determine which messages the method handles.

For example:

```
JMSType='car' AND color='blue' AND weight>2500
```

The selector string can refer to fields in the JMS message header and fields in the message properties. Message selectors cannot reference message body values.

Data type                      String, whose syntax is based on a subset of the SQL92 conditional expression syntax.

**Handle late responses:** Select this checkbox if you want to generate a method to handle late responses.

**Late response handler listener port name:** The name of the input port used to handle late responses.

This string must match the name of an input port defined in the Administrative Console.

Data type                      String

## Configuring deployment attributes for a sender bean

Use this task to configure the deployment attributes for a sender bean.

You can specify deployment attributes on each EJB method.

Changes to the deployment attributes override the values defined when the sender bean was developed and deployment code was generated for the application.

**Note:** After deployment code has been generated for an application, the deployable archive is renamed with the prefix `Deployed_`. Any subsequent changes to the archive from within the Application Assembly Tool are applied to the version of the archive that existed prior to code generation. To see changes reflected in your application, you must regenerate deployment code and re-install the deployable archive.

To change the deployment attributes for a sender bean, you can use the Application Assembly Tool to complete the following steps:

### Steps for this task

1. Launch the Application Assembly Tool.
2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, click **File-> Open** then select the EAR file.
3. In the navigation pane, select the sender bean instance; for example, expand *ejb\_module\_instance*-> **Extended messaging** then select the bean instance.  
A property dialog notebook for the sender bean is displayed in the property pane.
4. In the property pane, specify appropriate deployment attributes:

#### **Output port**

This is the name of the output port to be used to send messages.

#### **Handle late responses**

Select this checkbox if the sender bean is to handle late responses. If you select this checkbox, also specify the following properties:

**ReplyTimeout** and **Late response handler listener port name**.

**ReplyTimeout**

For a sender bean that has been developed to handle late responses, this is the time after which responses are considered late. This property is used if a response timeout is not specified on a sender method call.

**Late response handler listener port name**

For a sender bean that has been developed to handle late responses, this is the name of the listener port to be used for late responses.

For more information about the deployment attributes for sender beans, see Extended messaging assembly properties for enterprise beans.

5. To apply the changes and close the Application Assembly Tool, click **OK**. Otherwise, to apply the values but keep the property dialog open for additional edits, click **Apply**.
6. **(Optional)** To see changes reflected in your application, regenerate deployment code and reinstall the deployable archive.

---

## Configuring extended messaging service resources

Use these tasks with the WebSphere Administrative console to configure resources needed by the extended messaging service and applications that use extended messaging.

You can use WebSphere Application Server system management to configure resources needed by the extended messaging service and applications that use extended messaging.

For more information about the tasks involved, see the following topics:

- Adding a new input port
- Adding a new output port
- Configuring an input port
- Configuring an output port
- Configuring a listener port to handle late responses

### Adding a new input port

Use this task to add a new input port to WebSphere Application Server.

An input port is for use by an application that uses extended messaging.

During this task you configure the initial properties of the input port. You can later change the properties of the port, as described in Configuring an input port.

To add a new input port, complete the following steps:

#### Steps for this task

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Resources-> Extended messaging provider**  
This displays resources for extended messaging in the content pane.
3. In the Additional Properties table of the content pane, select **Input ports**  
This displays a list of the input ports in the content pane.
4. Click **New**.
5. Specify appropriate properties of the input port.



6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

## Adding a new output port

Use this task to add a new output port to WebSphere Application Server, and configure its properties, for use by an application that uses extended messaging.

During this task you configure the initial properties of the output port. You can later change the properties of the port, as described in [Configuring an output port](#).

To add a new output port, complete the following steps:

### Steps for this task

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Resources-> Extended messaging provider**  
This displays resources for extended messaging in the content pane.
3. In the Additional Properties table of the content pane, select **Output ports**  
This displays a list of the output ports in the content pane.
4. Click **New**.
5. Specify appropriate properties of the output port.
6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

## Configuring an input port

Use this task to change the properties of an input port for use by an application that uses extended messaging.

To change the properties of an input port, complete the following steps:

### Steps for this task

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Resources-> Extended messaging provider**  
This displays resources for extended messaging in the content pane.
3. In the Additional Properties table of the content pane, select **Input ports**  
This displays a list of the input ports in the content pane.
4. Select the input port that you want to change.
5. Specify appropriate properties of the input port.
6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

## Configuring an output port

Use this task to change the properties of an output port for use by an application that uses extended messaging.

To change the properties of an output port, complete the following steps:

### Steps for this task

1. Start the WebSphere Administrative console.
2. In the navigation pane, select **Resources-> Extended messaging provider**  
This displays resources for extended messaging in the content pane.
3. In the Additional Properties table of the content pane, select **Output ports**  
This displays a list of the output ports in the content pane.
4. Select the output port that you want to change.
5. Specify appropriate properties of the output port.
6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

## Extended messaging service settings

Use this page to enable or disable the extended messaging service.

The Extended Messaging Service provides run-time service for the support of extended messaging.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Extended Messaging Service** .

### Startup

Specifies whether the server will attempt to start the extended messaging service.

|         |  |
|---------|--|
| Default | Selected   |
| Range   | <b>Selected</b><br>When the application server starts, it attempts to start the extended messaging service automatically.  |
|         | <b>Cleared</b><br>The server does not try to start the extended messaging service. If extended messaging is to be used in applications that run on this server, the system administrator must start the extended messaging service manually or select this property then restart the server. |

### Late response handling extension collection

Use this page to view the configuration properties of late response handling extensions.

Late response handling extensions enable the handling of late responses with extended messaging

To view this administrative console page, click **Application Servers** > *server\_name* > **Extended Messaging Service** > **Listener Port Extensions** .

**Enabled:** Specifies whether the handling of late responses is enabled.

|       |  |
|-------|--|
| Range | <b>Selected</b><br>Handling of late responses is enabled.    |
|       | <b>Cleared</b><br>Handling of late responses is not enabled. |

**Request Interval:** Specifies the interval that elapses between checking for late responses.

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | milliseconds  |
| Default   | 5   |
| Range     | An integer number of milliseconds, greater than or equal to 0: <ul style="list-style-type: none"><li>• 0 indicates that the late response handler continually checks for requests</li><li>• Other values are an integer number of milliseconds between checks for requests.</li></ul> |

**Request Timeout:** Specifies the duration of time after which to give up waiting for a response.

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | seconds   |
| Default   | 0   |
| Range     | An integer number of milliseconds, greater than or equal to -1: <ul style="list-style-type: none"><li>• -1 indicates that requests to handle late responses are never discarded.</li><li>• Other values are an integer number of milliseconds after which requests are discarded.</li></ul> |

**Listener Ports:** Specifies the name of the listener port to be used to handle late responses.

### Late response handling extension settings

Use this page to configure late response handling extensions.

To view this administrative console page, click **Application Servers** > *server\_name* > **Extended Messaging Service** > **Listener Port Extensions** > *extension\_name* .

**Enabled:** Specifies whether the handling of late responses is enabled.

|       |  |
|-------|--|
| Range | <b>Selected</b><br>Handling of late responses is enabled.    |
|       | <b>Cleared</b><br>Handling of late responses is not enabled. |

**Request Interval:** Specifies the interval that elapses between checking for late responses.

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | milliseconds  |
| Default   | 5   |
| Range     | An integer number of milliseconds, greater than or equal to -1: <ul style="list-style-type: none"><li>• -1 indicates that requests to handle late responses are never discarded.</li><li>• Other values are an integer number of milliseconds after which requests are discarded.</li></ul> |

**Request Timeout:** Specifies the duration of time after which to give up waiting for a response.

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | seconds   |
| Default   | 0   |
| Range     | An integer number of milliseconds, greater than or equal to -1: <ul style="list-style-type: none"><li>• -1 indicates that requests to handle late responses are never discarded.</li><li>• Other values are an integer number of milliseconds after which requests are discarded.</li></ul> |

**Listener Ports:** Specifies the name of the listener port to be used to handle late responses.

## Extended messaging provider settings

Use this page to manage extended messaging providers.

The extended messaging provider manages resources defined for use with extended messaging.

To view this administrative console page, click **Resources > Extended Messaging Providers** .

### Name

The name of the resource provider.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Range     | 1 through 30 ASCII characters |

### Description

An optional description for the resource factory.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

### Input port collection

Use this page to view the configuration properties of input ports..

An input port specifies the properties needed by receiver beans as session beans. Receiver beans as message-driven beans do not need an input port, because the properties needed are associated with the deployed message-driven bean and the Message Listener service.

To view this administrative console page, click **Resources > Extended Messaging Providers > Input Port** .

**Name:** The name by which the input port is known for administrative purposes.

|           |                        |
|-----------|------------------------|
| Data type | String                 |
| Units     | En_US ASCII characters |

**JNDI Name:** The JNDI name for the resource.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** A description of the input port, for administrative purposes.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Category:** A string that can be used to classify or group the resource.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Range     | 1 through 30 ASCII characters |

**JMS Connection Factory JNDI Name:** The JNDI name for the JMS connection factory to be used by the input port; for example, `jms/connFactory1`.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**JMS Destination JNDI Name:** The JNDI name for the destination to be used by the input port; for example, `jms/destn1`.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**JMS Acknowledgement Mode:** JMS acknowledgment mode to be used for acknowledging messages.

This property applies only to message-driven beans that use bean-managed transaction demarcation (**Transaction type** is set to Bean).

|         |                  |
|---------|------------------|
| Default | Auto Acknowledge |
|---------|------------------|

Range

**Auto Acknowledge**

The session automatically acknowledges a message in either of the following cases:

- When the session has successfully returned from a call to receive a message.
- When the session has called a message listener to process the message and received a successful response from that listener.

**Dups OK Acknowledge**

The session acknowledges only the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

**Destination Type:** The type of the JMS resource.

Default

Queue

Range

**Queue** The receiver bean receives messages from a queue destination.

**Topic** The receiver bean receives messages from a topic destination.

**Subscription durability:** [Topic destinations only.] Specifies whether a JMS topic subscription is durable or non-durable.

Default

Durable

Range

**Durable**

A subscriber registers a durable subscription with a unique identity that is retained by JMS. Subsequent subscriber objects with the same identity resume the subscription in the state it was left in by the earlier subscriber. If there is no active subscriber for a durable subscription, JMS retains the subscription's messages until they are received by the subscription or until they expire.

**Nondurable**

Nondurable subscriptions last for the lifetime of their subscriber object. This means that a client sees the messages published on a topic only while its subscriber is active. If the subscriber is not active, the client is missing messages published on its topic.

**Reply JMS Connection Factory JNDI Name:** JNDI name of the JMS Connection Factory to be used for replies.

Data type

String

**Reply JMS Destination JNDI Name:** JNDI name of the JMS Destination to be used for replies.

Data type

String

## Input port settings

Use this page to configure an input port.

To view this administrative console page, click **Resources > Extended Messaging Providers > Input Port > *inputport\_name*** .

**Name:** The name by which the input port is known for administrative purposes.

|           |                        |
|-----------|------------------------|
| Data type | String                 |
| Units     | En_US ASCII characters |

**JNDI Name:** The JNDI name for the resource.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** A description of the input port, for administrative purposes.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Category:** A string that can be used to classify or group the resource.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Range     | 1 through 30 ASCII characters |

**JMS Connection Factory JNDI Name:** The JNDI name for the JMS connection factory to be used by the input port; for example, *jms/connFactory1*.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**JMS Destination JNDI Name:** The JNDI name for the destination to be used by the input port; for example, *jms/destn1*.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**JMS Acknowledgement Mode:** JMS acknowledgment mode to be used for acknowledging messages.

This property applies only to message-driven beans that use bean-managed transaction demarcation (**Transaction type** is set to Bean).

|         |                  |
|---------|------------------|
| Default | Auto Acknowledge |
|---------|------------------|

Range

**Auto Acknowledge**

The session automatically acknowledges a message in either of the following cases:

- When the session has successfully returned from a call to receive a message.
- When the session has called a message listener to process the message and received a successful response from that listener.

**Dups OK Acknowledge**

The session acknowledges only the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

**Destination Type:** The type of the JMS resource.

Default

Queue

Range

**Queue** The receiver bean receives messages from a queue destination.

**Topic** The receiver bean receives messages from a topic destination.

**Subscription durability:** [Topic destinations only.] Specifies whether a JMS topic subscription is durable or non-durable.

Default

Durable

Range

**Durable**

A subscriber registers a durable subscription with a unique identity that is retained by JMS. Subsequent subscriber objects with the same identity resume the subscription in the state it was left in by the earlier subscriber. If there is no active subscriber for a durable subscription, JMS retains the subscription's messages until they are received by the subscription or until they expire.

**Nondurable**

Nondurable subscriptions last for the lifetime of their subscriber object. This means that a client sees the messages published on a topic only while its subscriber is active. If the subscriber is not active, the client is missing messages published on its topic.

**Reply JMS Connection Factory JNDI Name:** JNDI name of the JMS Connection Factory to be used for replies.

Data type

String

**Reply JMS Destination JNDI Name:** JNDI name of the JMS Destination to be used for replies.

Data type

String



## Output port collection

Use this page to view the configuration properties of output ports.

The Output port defines the parameters required by the extended messaging sender bean. These properties define the destination for the message being sent, together with optional details if a response is expected.

To view this administrative console page, click **Resources > Extended Messaging Providers > Output Port** .

**Name:** The name by which the output port is known for administrative purposes.

Data type                      String

**JNDI Name:** The JNDI name for the output port.

Data type                      String

**Description:** A description of the output port, for administrative purposes.

Data type                      String

**Category:** A string that can be used to classify or group the resource.

Data type                      String

**JMS Connection factory JNDI name:** The JNDI name for the JMS connection factory to be used by the output port; for example, `jms/connFactory1`.

Data type                      String  
Units                              En\_US ASCII characters  
Range                              1 through 30 ASCII characters

**JMS Destination JNDI name:** The JNDI name for the destination to be used by the output port; for example, `jms/destn1`.

Data type                      String

**JMS Delivery Mode:** Specifies whether all messages sent to the destination are persistent or non-persistent.

Default                          Persistent  
Range                              **Persistent**  
    Messages put onto the destination are persistent.  
    **Nonpersistent**  
    Messages put onto the destination are not persistent.

**JMS Priority:** The message priority for this queue destination.

Data type                      Integer  
Default                          4  
Range                              0 to 9.

**JMS Time To Live:** The time in milliseconds after which messages on this queue expire.

|           |  |
|-----------|--|
| Data type | Integer  |
| Units     | Milliseconds   |
| Default   | 0  |
| Range     | 0 to <i>n</i>  |
|           | 0        messages never time out.                    |
|           | <i>n</i> messages time out in <i>n</i> milliseconds. |

**JMS Disable Message I.D.:** Specifies that the system should not generate a JMS message ID.

|         |   |
|---------|---|
| Default | Cleared   |
| Range   | <b>Selected</b><br>The system does not generate message IDs.      |
|         | <b>Cleared</b><br>The system generates message IDs automatically. |

**JMS Disabled Message Time Stamp:** Specifies that the system should not generate a JMS message timestamp.

|         |   |
|---------|---|
| Default | Cleared   |
| Range   | <b>Selected</b><br>Message time stamps are added automatically to messages sent.    |
|         | <b>Cleared</b><br>Message time stamps are not added automatically to messages sent. |

**Response JMS Connection Factory JNDI name:** The JNDI name for the JMS connection factory to be used for response messages handled by the output port; for example, `jms/connFactory1`.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Range     | 1 through 30 ASCII characters |

**Response JMS Destination JNDI name:** The JNDI name for the destination to be used for response messages handled by the output port; for example, `jms/destn1`.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

## Output port settings

Use this page to configure an output port.

An output port specifies the properties needed by sender beans to define the destination for the message being sent, and other optional properties if a response is expected. The output port is associated with the sender bean at deployment time.

To view this administrative console page, click **Resources > Extended Messaging Providers > Output Port > *outputport\_name*** .

**Name:** The name by which the output port is known for administrative purposes.

Data type                String

**JNDI Name:** The JNDI name for the output port.

Data type                String

**Description:** A description of the output port, for administrative purposes.

Data type                String

**Category:** A string that can be used to classify or group the resource.

Data type                String

**JMS Connection factory JNDI name:** The JNDI name for the JMS connection factory to be used by the output port; for example, *jms/connFactory1*.

Data type                String  
Units                    En\_US ASCII characters  
Range                    1 through 30 ASCII characters

**JMS Destination JNDI name:** The JNDI name for the destination to be used by the output port; for example, *jms/destn1*.

Data type                String

**JMS Delivery Mode:** Specifies whether all messages sent to the destination are persistent or non-persistent.

Default                  Persistent  
Range                    **Persistent**  
                              Messages put onto the destination are persistent.  
                              **Nonpersistent**  
                              Messages put onto the destination are not persistent.

**JMS Priority:** The message priority for this queue destination.

Data type                Integer  
Default                   4  
Range                    0 to 9.

**JMS Time To Live:** The time in milliseconds after which messages on this queue expire.

Data type                Integer  
Units                    Milliseconds

|         |  |
|---------|--|
| Default | 0  |
| Range   | 0 to <i>n</i>  |
|         | 0 messages never time out.                           |
|         | <i>n</i> messages time out in <i>n</i> milliseconds. |

**JMS Disable Message I.D.:** Specifies that the system should not generate a JMS message ID.

|         |   |
|---------|---|
| Default | Cleared   |
| Range   | <b>Selected</b><br>The system does not generate message IDs.      |
|         | <b>Cleared</b><br>The system generates message IDs automatically. |

**JMS Disabled Message Time Stamp:** Specifies that the system should not generate a JMS message timestamp.

|         |   |
|---------|---|
| Default | Cleared   |
| Range   | <b>Selected</b><br>Message time stamps are added automatically to messages sent.    |
|         | <b>Cleared</b><br>Message time stamps are not added automatically to messages sent. |

**Response JMS Connection Factory JNDI name:** The JNDI name for the JMS connection factory to be used for response messages handled by the output port; for example, `jms/connFactory1`.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Range     | 1 through 30 ASCII characters |

**Response JMS Destination JNDI name:** The JNDI name for the destination to be used for response messages handled by the output port; for example, `jms/destn1`.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

---

## Troubleshooting extended messaging

Use this overview task to help resolve a problem that you think is related to the extended messaging service.

The extended messaging service uses the standard WebSphere Application Server RAS facilities. If you encounter a problem that you think might be related to the extended messaging service, complete the following stages:

### Steps for this task

1. Check for extended messaging service messages in the application server's SystemOut log at `was_home\logs\server\SystemOut`.

Any error messages associated with the extended messaging service are labelled with **EMSG**. The error message indicates the nature of the problem and provides some detail. The extended messaging service issues EMSG error messages if it fails to initialize, parse its configuration file, or encounters some runtime error.

2. Check for more messages in the application server's SystemOut log.  
If the JMS server is running, but you have problems accessing JMS resources, check the SystemOut log file, which should contain more error messages and extra details about the problem.
3. Check the Release Notes for specific problems and workarounds  
The section *Possible Problems and Suggested Fixes* of the Release Notes, available from the WebSphere Application Server library web site, is updated regularly to contain information about known defects and their workarounds. Check the latest version of the Release Notes for any information about your problem. If the Release Notes does not contain any information about your problem, you can also search the Technotes database on the WebSphere Application Server web site.
4. Check for problems with the WebSphere Messaging functions or message-driven beans  
For more information about troubleshooting WebSphere Messaging, see the related topics listed at the bottom of this file.
5. **(Optional)** Get a detailed exception dump for extended messaging.  
If the information obtained in the preceding steps is still inconclusive, you can enable the application server debug trace for the "Messaging" group to provide a detailed exception dump.

---

## Extended Messaging: Resources for learning


Use the following links to find relevant supplemental information about Extended Messaging. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.


View links to additional information about:

- Programming model and decisions
- Programming specifications
- Other





### Programming model and decisions

-  Sun's Java Message Service (JMS) specification documentation

### Programming specifications

-  Java Message Service API, 1.0.2
-  Enterprise JavaBeans Technology Downloads & Specifications

## Other

-  WebSphere Application Server Enterprise Version 5 Overview: Extended J2EE Development Accelerators
-  Listing of PDF files to learn about WebSphere Application Server Version 5
-  Listing of all IBM WebSphere Application Server Redbooks
-  Listing of all IBM WebSphere Application Server Whitepapers
-  WebSphere Application Server Enterprise Edition 4.0: A Programmer's Guide

---

## Chapter 6. Using message-driven beans in applications

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface.

Message-driven beans (a type of enterprise bean defined in the EJB 2.0 specification) extend the base JMS support and the Enterprise JavaBean component model to provide automatic asynchronous messaging. When a message arrives on a destination, a listener passes the message to a new instance of a user-developed message-driven bean for processing.

You can use WebSphere Studio Application Developer to develop applications that use message-driven beans. You can use the WebSphere Application Server runtime tools, like the administrative console, to deploy and administer applications that use message-driven beans.

For more information about implementing WebSphere enterprise applications that use message-driven beans, see the following topics:

- An overview of message-driven beans
- Designing an enterprise application to use a message-driven bean
- Developing an enterprise application to use a message-driven bean
- Deploying an enterprise application to use a message-driven bean
- Configuring message listener resources for message-driven beans
- Troubleshooting problems with message-driven beans

---

### Message-driven beans - an overview

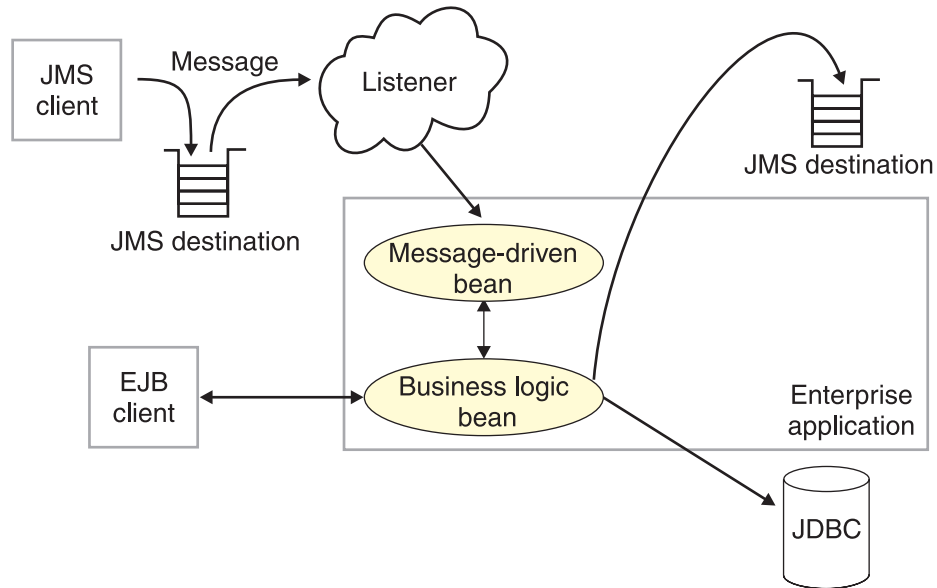
WebSphere Application Server supports automatic asynchronous messaging with *message-driven beans* (a type of enterprise bean defined in the EJB 2.0 specification). Messaging with message-driven beans is shown in the figure Message-driven beans and the message listener service.

The support for message-driven beans is based on the message listener service, which comprises a *listener manager* that controls and monitors one or more *listeners*. Each listener monitors a JMS destination for incoming messages. When a message arrives on the destination, the listener passes the message to a new instance of a user-developed message-driven bean (an enterprise bean) for processing. The listener then looks for the next message without waiting for the bean to return.

Messages arriving at a destination being processed by a listener have no client credentials associated with them; the messages are anonymous. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component. For more information about EJB security, see Enterprise bean component security in the WebSphere Application Server version 5 Security PDF document or information center.

You are recommended to develop a message-driven bean to delegate the business processing of incoming messages to another enterprise bean, to provide clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client.

**Message-driven beans and the message listener service.** This figure shows an incoming message being passed by a JMS listener to a message-driven bean, which passes the message on to a business logic bean for business processing. This messaging is controlled by the listener manager. For more information, see the text that accompanies this figure.

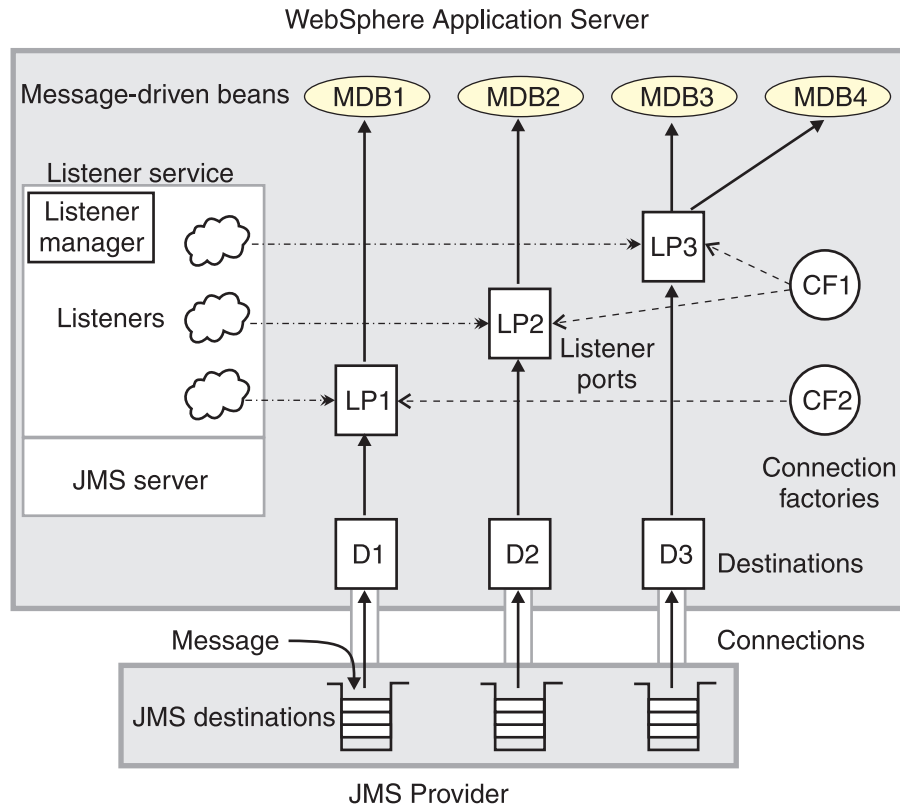


## Message-driven beans - components

The WebSphere Application Server support for message-driven beans is based on JMS message listeners and the message listener service, and builds on the base support for JMS. The main components of WebSphere Application Server support for message-driven beans are shown in the following figure and described after the figure:

**The main components for message-driven beans.** This figure shows the main components of WebSphere support for message-driven beans, from JMS provider through a connection to a destination, listener port, then deployed message-driven bean that processes the message retrieved from the destination. Each listener port defines the association between a connection factory, destination, and a deployed message-driven bean. The other main components are the message listener service, which comprises a listener for each listener port, all controlled by the same listener manager. For more information, see the text that accompanies this figure.





The *message listener service* is an extension to the JMS functions of the JMS provider and provides a *listener manager*, which controls and monitors one or more JMS *listeners*.

Each listener monitors either a JMS queue destination (for point-to-point messaging) or a JMS topic destination (for publish/subscribe messaging).

A *connection factory* is used to create connections with the JMS provider for a specific JMS queue or topic destination. Each connection factory encapsulates the configuration parameters needed to create a connection to a JMS destination.

A *listener port* defines the association between a connection factory, a destination, and a deployed *message-driven bean*. Listener ports are used to simplify the administration of the associations between these resources.

When a deployed message-driven bean is installed, it is associated with a listener port and the listener for a destination. When a message arrives on the destination, the listener passes the message to a new instance of a message-driven bean for processing.

When an application server is started, it initializes the listener manager based on the configuration data. The listener manager creates a dynamic session thread pool for use by listeners, creates and starts listeners, and during server termination controls the cleanup of listener message service resources. Each listener completes several steps for the JMS destination that it is to monitor, including:

- Creating a JMS server session pool, and allocating JMS server sessions and session threads for incoming messages.

- Interfacing with JMS ASF to create JMS connection consumers to listen for incoming messages.
- If specified, starting a transaction and requesting that it is committed (or rolled back) when the EJB method has completed.
- Processing incoming messages by invoking the onMessage() method of the specified enterprise bean.

## Message-driven beans - transaction support

Message-driven beans can handle messages read from JMS destinations within the scope of a transaction. If transaction handling is specified for a JMS destination, the JMS listener starts a global transaction *before* it reads any incoming message from that destination. When the message-driven bean processing has finished, the JMS listener commits or rolls back the transaction (using JTA transaction control).

### Note:

- All messages retrieved from a specific destination have the same transactional behavior.

If messages are queued to be sent within a global transaction they are sent when the transaction is committed. If the processing of a message causes the transaction to be rolled back, then the message that caused the bean instance to be invoked is left on the JMS destination.

You can configure the **Maximum retries** property of the listener port to define the maximum number of times the listener attempts to read a message from a destination. When the Max retries limit is reached, the listener for that destination is stopped. When you have resolved the problem, you must then restart the listener.

---

## Designing an enterprise application to use message-driven beans

This topic describes things to consider when designing an enterprise application to use message-driven beans.

The considerations in this topic are based on a generic enterprise application that uses one message-driven bean to retrieve messages from a JMS queue destination and passes the messages on to another enterprise bean that implements the business logic.

To design an enterprise application to use message-driven beans, complete the following steps:

### Steps for this task

1. Identify the JMS resources that the application is to use.

This helps to identify the properties of resources that need to be used within the application and configured as application deployment descriptors or within WebSphere Application Server.

| JMS resource type        | Properties  |
|--------------------------|---|
| Queue connection factory | <b>Name:</b> SamplePtoPQueueConnectionFactory<br><b>JNDI Name:</b> Sample/JMS/QCF |
| Queue destination        | <b>Name:</b> Q1<br><b>JNDI Name:</b> Sample/JMS/Q1                                |

| JMS resource type                           | Properties  |
|---|---|
| Listener port (for the destination)         | <b>Name:</b> SamplePtoPListenerPort<br><b>Connection Factory JNDI Name:</b> Sample/JMS/QCF<br><b>Destination JNDI Name:</b> Sample/JMS/Q1<br><b>Maximum Sessions:</b> 5<br><b>Maximum Retries:</b> 10<br><b>Maximum Messages:</b> 1 |
| Message-driven bean (deployment properties) | <b>Name:</b> JMSppSampleMDBBean<br><b>Transaction type:</b> Container<br><b>Destination type:</b> Queue<br><b>Listener port name:</b> SamplePtoPListenerPort  |
| Business logic bean                         | <b>Name:</b> MyLogicBean  |

Ensure that you use consistent values where needed; for example, the JNDI names for the connection factory and destination must be the same for both those resources and the equivalent properties of the listener port.

## 2. Separation of business logic.

You are recommended to develop a message-driven bean to delegate the business processing of incoming messages to another enterprise bean. This provides clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client.

## 3. Security considerations.

Messages arriving at a destination being processed by a listener have no client credentials associated with them; the messages are anonymous. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component. For more information about EJB security, see Enterprise bean component security in the WebSphere Application Server version 5 Security PDF document or information center.

## 4. General JMS considerations

For Publish/Subscribe messaging, choose the JMS server port to be used depending on your needs for transactions or performance:

### Queued port

The TCP/IP port number of the listener port used for all point-to-point and Publish/Subscribe support.

### Direct port

The TCP/IP port number of the listener port used for direct TCP/IP connection (non-transactional, non-persistent, and non-durable subscriptions only) for Publish/Subscribe support.

**Note:** Message-driven beans cannot use the direct listener port for Publish/Subscribe support. Therefore, any topic connection factory configured with **Portset to Direct** cannot be used with message-driven beans.

A non-durable subscriber can only be used in the same transactional context (for example, a global transaction or an unspecified transaction context) that existed when the subscriber was created. For more information about this context restriction, see The effect of transaction context on non-durable subscribers in the WebSphere Application Server version 5 Resources PDF document or information center.

---

## Developing an enterprise application to use message-driven beans

Use this task to develop an enterprise application to use a message-driven bean. The message-driven bean is invoked by a JMS listener when a message arrives on the input queue that the listener is monitoring.

You are recommended to develop the message-driven bean to delegate the business processing of incoming messages to another enterprise bean, to provide clear separation of message handling and business processing. This also enables the business processing to be invoked by either the arrival of incoming messages or, for example, from a WebSphere J2EE client. Responses can be handled by another enterprise bean acting as a sender bean, or handled in the message-driven bean.

You develop an enterprise application to use a message-driven bean like any other enterprise bean, except that a message-driven bean does not have a home interface or a remote interface.

This topic describes how to develop a completely new message-driven bean class. If you have a WAS 4.0 enterprise application that uses the JMS listener, you can migrate that application to use message-driven beans, as described in *Migrating a WAS 4.0 JMS listener application to use message-driven beans*.

For more information about writing the message-driven bean class, see *Creating a message-driven bean* in the WebSphere Studio help bookshelf.

To develop an enterprise application to use a message-driven bean, complete the following steps:

### Steps for this task

1. Creating the Enterprise Application project, as described in the WebSphere Studio article .
2. Creating the message-driven bean class.

You can use the New Enterprise Bean wizard of WebSphere Studio Application Developer to create an enterprise bean with a bean type of Message-driven bean. The wizard creates appropriate methods for the type of bean.

By convention, the message bean class is named *nameBean*, where *name* is the name you assign to the message bean; for example:

```
public class MyJMSppMDBBean implements MessageDrivenBean, MessageListener
```

The message-driven bean class must define and implement the following methods:

- `onMessage(message)`, which must meet the following requirements:
  - The method must have a single argument of type `javax.jms.Message`.
  - The throws clause must *not* define any application exceptions.
  - If the message-driven bean is configured to use bean-managed transactions, it must call the `javax.transaction.UserTransaction` interface to scope the transactions. Because these calls occur inside the `onMessage()` method, the transaction scope does not include the initial message receipt. This means the application server is given one attempt to process the message.

To handle the message within the `onMessage()` method (for example, to pass the message on to another enterprise bean), you use standard JMS. (This is known as bean-managed messaging.)

- `ejbCreate()`

You must define and implement an `ejbCreate` method for each way in which you want a new instance of an enterprise bean to be created.

- `ejbRemove()`.

This method is invoked by the container when a client invokes the `remove` method inherited by the enterprise bean's home interface from the `javax.ejb.EJBHome` interface. This method must contain any code that you want to execute before an enterprise bean instance is removed from the container (and the associated data is removed from the data source).

For example, the following code extract shows how to access the text and the JMS MessageID, from a JMS message of type `TextMessage`:

**Code example: The `onMessage()` method of a message bean.** This figure shows a code extract for a basic `onMessage()` method of a sample message-driven bean. The method unpacks the incoming text message to extract the text and message identifier and calls a private `putMessage` method (defined within the same message bean class) to put the message onto another queue.

```
public void onMessage(javax.jms.Message msg)
{
    String text      = null;
    String messageID = null;

    try
    {
        text = ((TextMessage)msg).getText();

        System.out.println("senderBean.onMessage(), msg text2: "+text);

        //
        // store the message id to use as the Correlator value
        //
        messageID = msg.getJMSMessageID();

        // Call a private method to put the message onto another queue
        putMessage(messageID, text);
    }
    catch (Exception err)
    {
        err.printStackTrace();
    }
    return;
}
```

The result of this step is a message-driven bean that can be assembled into an `.EAR` file for deployment.

3. Assembling and packaging the application for deployment.

You can use WebSphere Studio to assemble and package the application for deployment.

### Results

The result of this task is an `.EAR` file, containing an application message-driven bean, that can be deployed in WebSphere Application Server.

## What to do next

After you have developed an enterprise application to use message-driven beans, configure and deploy the application; for example, define the listener ports for the message-driven beans and, optionally, change the deployment descriptor attributes for the application. For more information about configuring and deploying an application that uses message-driven beans, see "Deploying an enterprise application to use message-driven beans"

## **Migrating a JMS listener application to use message-driven beans**

Use this task to migrate an enterprise application that uses message beans with the JMS Listener from WebSphere Application Server 4.0 to use EJB 2.0 message-driven beans.

This task uses a command line utility, `mb2mdb`, that takes as its input either a deployed `MessageBean.jar` module or a deployed Enterprise Application (`.ear`) that contains a message bean, along with the JMS listener configuration XML file that defines the WebSphere Application Server 4.0 message beans. The result is a new `.jar/.ear` module that can then be deployed directly into a WebSphere Application Server 5.0 application server.

You can display the usage help for the migration utility, by typing the command `mb2mdb` at a command line.

To migrate a WebSphere Application Server 4.0 enterprise application that uses message beans to use EJB 2.0 message-driven beans, type the following command at an operating system command line:

```
mb2mdb inputMB.jar-ear jmsListenerConfig.xml workingDirectory outputMDB.jar-ear options
```

Where:

### *inputMB.jar-ear*

The name of the deployed WebSphere Application Server 4.0 jar or ear file containing a stateless session message bean.

### *jmsListenerConfig.xml*

The name of the XML configuration file used to configure the WebSphere Application Server 4.0 JMS listeners.

### *workingDirectory*

The name of a new or existing directory that is used to generate the new message-driven bean and package the `outputMDB.jar` or `.ear` file.

**Note:** By default, the tool clears the working directory after it has completed. If you want to preserve the contents of the working directory, you must specify the `-keep` option.

### *outputMDB.jar-ear*

The name of the output `.jar` or `.ear` file for the migrated message-driven bean application.

### *options*

An optional set of parameters that you can use to control the `mb2mdb` utility.

**-keep** This prevents the tool from clearing out the working directory after completion.

**-verbose**

This causes the tool to display informational messages as to the progress of the migration and its parameters.

**-map listenerHome=bindingHome**

This option provides a mechanism to map between the JNDIHomeName specified for a listener in the JMS listener configuration XML file and the default binding home name specified in the *inputMB.jar-ear* file.

If the *jmsListenerConfig.xml* file contains a deployed EJB home JNDI name that is different to the default binding within the *inputMB.jar-ear*, use this option to map between the two names.

This enables you to install the output .jar or .ear file for the message-driven bean into an application server and bind the bean with a different JNDIHomeName than is specified in the bean's bindings.xmi.

### Results

The result of this task is a new .jar or .ear file for a message-driven bean that can then be deployed directly into a WebSphere Application Server 5.0 application server.

To successfully install the .jar or .ear file, you need to bind the message-driven bean against a listener port defined to the message listener service of the application server. You need to have used the WebSphere Application Server administrative console to define the listener port, which defines the JMS connection factory and destination that a message-driven bean bound to it listens on. For more information about installing and configuring a .jar or .ear file for a message-driven bean, see *Deploying an enterprise application to use message-driven beans*.

---

## Deploying an enterprise application to use message-driven beans

Use this task to deploy an enterprise application to use message-driven beans.

This task description assumes that you have an .EAR file, which contains an application enterprise bean with code for message-driven beans, that can be deployed in WebSphere Application Server.

To deploy an enterprise application to use message-driven beans, complete the following steps:

### Steps for this task

1. Use the WebSphere administrative console to define the listener ports for the application, as described in *Adding a new listener port*.
2. For each message-driven bean in the application, use the application assembly tool to configure the deployment attributes to match the listener port definitions, as described in *"Configuring deployment attributes for a message-driven bean"*.
3. Use the WebSphere administrative console to install the application  
This stage is a standard WebSphere Application Server task, as described in *"Installing a new application"*.

When you install the application, you are prompted to specify the name of the listener port that the application is to use for late responses. Select the listener port, then click **OK**.

## Configuring deployment attributes for a message-driven bean

Use this task to configure the message-driven beans deployment attributes for an enterprise bean, to override the deployment attributes defined within the application EAR file.

This task description assumes that you have an EAR file, which contains an application enterprise bean developed as a message-driven bean, that can be deployed in WebSphere Application Server.

**Note:** After deployment code has been generated for an application, the deployable archive is renamed with the prefix `Deployed_`. Any subsequent changes to the archive from within the Application Assembly Tool are applied to the version of the archive that existed prior to code generation. To see changes reflected in your application, you must regenerate deployment code and re-install the deployable archive.

To configure the message-driven beans deployment attributes for an enterprise bean, use the the application assembly tool to configure the deployment attributes of the application to match the listener port definitions:

### Steps for this task

1. Launch the Application Assembly Tool.
2. Create or edit the application EAR file.

For example, to change attributes of an existing application, click **File-> Open** then select the the EAR file.

3. In the navigation pane, select the message-driven bean instance; for example, expand `ejb_module_instance-> Message-driven beans` then select the bean instance.

A property dialog notebook for the message-driven bean is displayed in the property pane.

4. Specify general deployment properties.
  - a. In the property pane, select the General tab.
  - b. Specify the following properties:

#### **Transaction type**

Whether the message bean manages its own transactions or the container manages transactions on behalf of the bean. All messages retrieved from a specific destination have the same transactional behavior. To enable the transactional behavior that you want, you must configure the JMS destination with the same transactional behavior as you configure for the message bean.

**Bean** The message bean manages its own transactions

#### **Container**

The container manages transactions on behalf of the bean

5. Specify advanced deployment properties.
  - a. In the property pane, select the Advanced tab.
  - b. Specify the following properties:



### **Message selector**

The JMS message selector to be used to determine which messages the message bean receives; for example:

```
JMSType='car' AND color='blue' AND weight>2500
```

The selector string can refer to fields in the JMS message header and fields in the message properties. Message selectors cannot reference message body values.

### **Acknowledge mode**

How the session acknowledges any messages it receives.

This property applies only to message-driven beans that uses bean-managed transaction demarcation (**Transaction type** is set to Bean).

#### **Auto Acknowledge**

The session automatically acknowledges a message when it has either successfully returned from a call to receive, or the message listener it has called to process the message successfully returns.

#### **Dups OK Acknowledge**

The session lazily acknowledges the delivery of messages. This is likely to result in the delivery of some duplicate messages if JMS fails, so it should be used only by consumers that are tolerant of duplicate messages.

As defined in the EJB specification, clients cannot use using Message.acknowledge() to acknowledge messages. If a value of CLIENT\_ACKNOWLEDGE is passed on the createxxxSession call, then messages are automatically acknowledged by the application server and Message.acknowledge() is not used.

### **Destination type**

Whether the message bean uses a queue or topic destination.

#### **Queue**

The message bean uses a queue destination.

**Topic** The message bean uses a topic destination.

### **Subscription durability**

Whether a JMS topic subscription is durable or non-durable.

#### **Durable**

A subscriber registers a durable subscription with a unique identity that is retained by JMS. Subsequent subscriber objects with the same identity resume the subscription in the state it was left in by the earlier subscriber. If there is no active subscriber for a durable subscription, JMS retains the subscription's messages until they are received by the subscription or until they expire.

#### **Nondurable**

Non-durable subscriptions last for the lifetime of their subscriber object. This means that a client sees the messages

published on a topic only while its subscriber is active. If the subscriber is not active, the client is missing messages published on its topic.

A non-durable subscriber can only be used in the same transactional context (for example, a global transaction or an unspecified transaction context) that existed when the subscriber was created. For more information about this context restriction, see *The effect of transaction context on non-durable subscribers in the WebSphere Application Server version 5 Resources PDF document* or information center.

6. Specify bindings deployment properties.
  - a. In the property pane, select the Bindings tab.
  - b. Specify the following property:

**Listener port name**

The name of the listener port for this message-driven bean.

7. To apply the changes and close the Application Assembly Tool, click **OK**. Otherwise, to apply the values but keep the property dialog open for additional edits, click **Apply**.
8. **(Optional)** To see changes reflected in your application, regenerate deployment code and reinstall the deployable archive.

---

## Configuring message listener resources for message-driven beans

Use the following tasks to configure resources needed by the message listener service to support message-driven beans.

- Configuring the message listener service
- Adding a new listener port
- Configuring a listener port
- Configuring security for message-driven beans

### Configuring the message listener service

Use this task to configure the properties of the message listener service for an application server.

To configure the properties of the message listener service for an application server, use the administrative console to complete the following steps:

Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**  
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.  
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**  
This displays the Message Listener Service properties in the content pane.
4. Specify appropriate properties of the message listener service.
5. Click **OK**.
6. Save your configuration.

7. **(Optional)** To have the changed configuration take effect, stop then restart the Application Server.

### Message listener service

The message listener service is an extension to the JMS functions of the JMS provider. It provides a listener manager that controls and monitors one or more JMS listeners, which each monitor a JMS destination on behalf of a deployed message-driven bean.

This panel displays links to the Additional Properties pages for Listener Ports, Thread Pool, and Custom Properties for the message listener service.

To view this administrative console page, click **Servers-> application\_server-> Message Listener Service**

**Name:** The name by which the message listener service is known for administrative purposes.

|           |             |
|-----------|-------------|
| Data type | String      |
| Default   | MsgLService |

**Description:** A description of the message listener service, for administrative purposes

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Thread pool:** Controls the maximum number of threads the Message Listener Service is allowed to run. Select this link to display the service thread pool properties.

Adjust this parameter when multiple message-driven beans are deployed in the same application server and the sum of their maximum session values exceeds the default value of 10.

|             |   |
|-------------|---|
| Data type   | Integer   |
| Units       | Not applicable  |
| Default     | Minimum: 10, maximum 50   |
| Range       | Not applicable  |
| Recommended | Set the minimum to the sum of all message-driven beans maximum session values. Set the maximum to anything equal or greater than the minimum. |

**Custom Properties:** Custom properties of the message listener service. The JMS server has two modes of operation, Application Server Facilities (ASF) and non-Application Server Facilities. ASF is meant to provide concurrency and transactional support for applications. Non-ASF bypasses that support to streamline the path length.

Use Non-ASF if:

- Message order is a strict requirement
- Concurrent PTP messages are desired
- The property `non.asf.receive.timeout` exists and has a value greater than 0

Do not use Non-ASF if concurrent publications and subscriptions messages are desired because ASF provides better throughput.

|             |  |
|-------------|--|
| Data type   | Integer  |
| Units       | Milliseconds   |
| Default     | ASF mode (custom property not created)   |
| Range       |  |
| Recommended | The value is the number of milliseconds it takes for a message to be delivered. If a timeout occurs, it must recycle causing extra work. This should be set to lower than the transaction timeout, but close to it with several (10 or more) seconds to spare. The seconds to spare should be more if under extreme loads in which threads are waiting long periods of time to get CPU cycles. |

**MQJMS.POOLING.TIMEOUT:** The number of milliseconds after which a connection in the pool is destroyed if it has not been used.

An MQSimpleConnectionManager allocates connections on a most-recently-used basis, and destroys connections on a least-recently-used basis. By default, a connection is destroyed if it has not been used for five minutes.

|           |              |
|-----------|--------------|
| Data type | Integer      |
| Units     | Milliseconds |
| Default   | 5 minutes    |
| Range     |              |

**MQJMS.POOLING.THRESHOLD:** The maximum number of unused connections in the pool.

An MQSimpleConnectionManager allocates connections on a most-recently-used basis, and destroys connections on a least-recently-used basis. By default, a connection is destroyed if there are more than ten unused connections in the pool.

|           |                       |
|-----------|-----------------------|
| Data type | Integer               |
| Units     | Number of connections |
| Default   | 10                    |
| Range     |                       |

**MAX.RECOVERY.RETRIES:** The maximum number of times that the listener service tries to get a message from a listener port before the associated listener is stopped, in the range 0 through 2147483647.

|           |                                   |
|-----------|-----------------------------------|
| Data type | Integer                           |
| Units     | Retry attempts                    |
| Default   | 0 (no retries)                    |
| Range     | 0 (no retries) through 2147483647 |

**RECOVERY.RETRY.INTERVAL:** The time in seconds between retry attempts by the listener service to get a message from a listener port.

|           |                      |
|-----------|----------------------|
| Data type | Integer              |
| Units     | Seconds              |
| Default   | 10                   |
| Range     | 1 through 2147483647 |

## Message listener port collection

The message listener ports configured in the administrative domain

This panel displays a list of the message listener ports configured in the administrative domain. Each listener port is used with a message-driven bean to automatically receive messages from an associated JMS destination. You can use this panel to add new listener ports or to change the properties of existing listener ports. For more information about the property fields for listener ports, see Listener port properties.

To view this administrative console page, click **Servers-> application\_server-> Message Listener Service-> Listener Ports**

### Listener port settings

A listener port is used to simplify administration of the association between a connection factory, destination, and deployed message-driven bean.

Use this panel to view or change the configuration properties of the selected listener port.

To view this administrative console page, click **Servers-> application\_server-> Message Listener Service-> Listener Ports-> listener\_port**

**Initial state:** The state that you want the listener port to have when the application server is next restarted

|           |   |
|-----------|---|
| Data type | Enum  |
| Units     | Not applicable  |
| Default   | Started   |
| Range     | <b>Started</b> When the application server is next started, the listener port is started automatically.<br><b>Stopped</b> When the application server is next started, the listener port is not started automatically. If message-driven beans are to use this listener port on the application server, the system administrator must start the port manually or select the Started value of this property then restart the application server. |

**Description:** A description of the listener port, for administrative purposes within IBM WebSphere Application Server.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Listener port:** The name by which the listener port is known for administrative purposes.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Connection factory JNDI name:** The JNDI name for the JMS connection factory to be used by the listener port; for example, jms/connFactory1.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Destination JNDI name:** The JNDI name for the destination to be used by the listener port; for example, `.jms/destn1`.

If the extended messaging service is to use this listener port to handle late responses, the value of this property must match the JMS response destination on the output port used by the sender bean. Also cannot use a temporary destination for late responses.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Maximum sessions:** Specifies the maximum number of concurrent JMS server sessions used by a listener to process messages.

Each session corresponds to a single listener thread, but does not control the number of concurrently processed messages. Adjust this parameter when the machine running the process application does not realize the available capacity and produces less throughput by running long processes.

This parameter helps to add some level of concurrency so that messages are received concurrently before the synchronized part of receiving the message occurs. A 40% increase in throughput is observed in a process application with long running processes on a NetFinity 5500 500 MHz, 4-way, 4GB RAM system.

An alternative to making this change in the administrative console is to change the maximum sessions value in the `server.xml` file under the listener ports stanza as follows:

```
<listenerPorts xmi:id="ListenerPort_1" name="bpeIntListenerPort"
description="Internal Listener Port for Process Choreographer"
connectionFactoryJNDIName=".jms/bpeCF" destinationJNDIName=".jms/bpeIntQueue"
maxSessions="5" maxRetries="10" maxMessages="1">
<stateManagement xmi:id="StateManageable_5" initialState="START" />
</listenerPorts>
```

|             |   |
|-------------|---|
| Data type   | Integer   |
| Units       | Sessions  |
| Default     | 1   |
| Range       | 1 through 2147483647  |
| Recommended | If message concurrency, that is multiple messages processed simultaneously, is desired, set the value to 2-4 sessions per system processor. Keep the value as low as possible to eliminate client thrashing. If a strict message order is desired, set the value to 4 so there is always a thread waiting in a hot state, blocked on receiving the message. |

**Maximum retries:** The maximum number of times that the listener tries to deliver a message before the listener is stopped, in the range 0 through 2147483647.

The maximum number of times that the listener tries to deliver a message to a message-driven bean instance before the listener is stopped.

|           |         |
|-----------|---------|
| Data type | Integer |
|-----------|---------|

|         |                                   |
|---------|-----------------------------------|
| Units   | Retry attempts                    |
| Default | 0 (no retries)                    |
| Range   | 0 (no retries) through 2147483647 |

**Maximum messages:** The maximum number of messages that the listener can process in one JMS server session.

Adjust this parameter when message order is not a strict requirement. This parameter controls the concurrency level in the listener service.

|             |  |
|-------------|--|
| Data type   | Integer  |
| Units       | Number of messages   |
| Default     | 1  |
| Range       | 1 through 2147483647   |
| Recommended | If message concurrency, that is multiple messages processed simultaneously, is desired, set the value to 2-4 sessions per system processor. Keep the value as low as possible to eliminate client thrashing. If a strict message order is desired, set the value to 1. |

## Adding a new listener port

Use this task to add a new listener port to the message listener service, so that message-driven beans can be associated with the port to retrieve messages.

To add a new listener port, use the administrative console to complete the following steps:

### Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**  
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.  
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**  
This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**.  
This displays a list of the listener ports.
5. In the content pane, click **New**.
6. Specify appropriate properties for the listener port.
7. Click **OK**.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.
9. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

### Results

If enabled, the listener port is started automatically when a message-driven bean associated with that port is installed.

## Configuring a listener port

Use this task to change the properties of an existing listener port, used by message-driven beans associated with the port to retrieve messages.

To configure the properties of a listener port, use the administrative console to complete the following steps:

### Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**  
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.  
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**  
This displays the Message Listener Service properties in the content pane.
4. In the content pane, click **Listener Ports**.  
This displays a list of the listener ports.
5. Click the listener port that you want to modify.  
This displays the properties of the listener port in the content pane.
6. Specify appropriate properties for the listener port.
7. Click **OK**.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.
9. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

## Deleting a listener port

Use this task to delete a listener port from the message listener service, to prevent message-driven beans associated with the port from retrieving messages.

To delete a listener port, use the administrative console to complete the following steps:

### Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**  
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.  
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**  
This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**.  
This displays a list of the listener ports.
5. In the content pane, select the checkbox for the listener port that you want to delete.
6. Click **Delete**. This action stops the port (needed to allow the port to be deleted) then deletes the port.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.



8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

## Configuring security for message-driven beans

Use this task to configure resource security and security permissions for message-driven beans.

Messages arriving at a listener port have no client credentials associated with them. The messages are anonymous.

To call secure enterprise beans from a message-driven bean, the message-driven bean needs to be configured with a RunAs Identity deployment descriptor. Security depends on the role specified by the RunAs Identity for the message-driven bean as an EJB component.

For more information about EJB security, see Enterprise bean component security in the WebSphere Application Server version 5 Security PDF document or information center. For more information about configuring security for your application, see Assembling secured applications.

JMS connections used by message-driven beans can benefit from the added security of using J2C container-managed authentication. To enable the use of J2C container authentication aliases and mapping, define a J2C container-managed alias on the JMS connection factory definition that the MDB is using to listen upon (defined by the **Connection factory JNDI name** property of the listener port). If defined, the listener uses the container-managed authentication alias for its JMSConnection security credentials instead of any application-managed alias. To set the container-managed alias, use the administrative console to complete the following steps:

### Steps for this task

1. To display the listener port settings, click **Servers-> application\_server-> Message Listener Service-> Listener Ports-> listener\_port**
2. To get the name of the JMS connection factory, look at the **Connection factory JNDI name** property.
3. Display the JMS connection factory properties. For example, to display the properties of a queue connection factory provided by the embedded WebSphere JMS provider, click **Resources-> WebSphere JMS Provider-> (In content pane, under Additional Properties) WebSphere Queue Connection Factories-> connection\_factory**
- 4.
5. Set the **Container-managed Authentication Alias** property.

## Administering listener ports

Use the following tasks to administer listener ports, which each define the association between a connection factory, a destination, and a message-driven bean.

You can use the WebSphere administrative console to administer listener ports, as described in the following tasks.

- Adding a new listener port

Use this task to create a new listener port, to specify a new association between a connection factory, a destination, and a message-driven bean. This enables deployed message-driven beans associated with the port to retrieve messages from the destination.

- **Configuring a listener port**  
Use this task to view or change the configuration properties of a listener port.
- **Starting a listener port**  
Use this task to start a listener port manually.
- **Stopping a listener port**  
Use this task to stop a listener port manually.

**Note:** If configured as enabled, a listener port is started automatically when a message-driven bean associated with that port is installed. You do not normally need to start or stop a listener port manually.

### Starting a listener port

Use this task to start a listener port on an application server, to enable the listeners for message-driven beans associated with the port to retrieve messages.

A listener is active, that is able to receive messages from a destination, if the deployed message-driven bean, listener port, and message listener service are all started. Although you can start these components in any order, they must all be in a started state before the listener can retrieve messages.

If configured as enabled, a listener port is started automatically when a message-driven bean associated with that port is installed. However, you can start a listener port manually, as described in this topic.

When a listener port is started, the listener manager tries to start the listeners for each message-driven bean associated with the port. If a message-driven bean is stopped, the port is started but the listener is not started, and remains stopped. If you start a message-driven bean, the related listener is started.

To start a listener port on an application server, use the administrative console to complete the following steps:

#### Steps for this task

1. **(Optional)** If you want the listener for a deployed message-driven bean to be able to receive messages at the port, check that the message-driven bean has been started.
2. In the navigation pane, select **Servers-> Application Servers**  
This displays a table of the application servers in the administrative domain.
3. In the content pane, click the name of the application server.  
This displays the properties of the application server in the content pane.
4. In the Additional Properties table, select **Message Listener Service**  
This displays the Message Listener Service properties in the content pane.
5. In the content pane, select **Listener Ports**.  
This displays a list of the listener ports.
6. **(Optional)** Select the checkbox for the listener port that you want to start.
7. Click **Start**.
8. To save your configuration, click **Save** on the task bar of the Administrative console window.

## Stopping a listener port

Use this task to stop a listener port on an application server, to prevent the listeners for message-driven beans associated with the port from retrieving messages.

When you stop a listener port as described in this topic, the listener manager stops the listeners for all message-driven beans associated with the port.

To stop a listener port on an application server, use the administrative console to complete the following steps:

### Steps for this task

1. In the navigation pane, select **Servers-> Application Servers**  
This displays a table of the application servers in the administrative domain.
2. In the content pane, click the name of the application server.  
This displays the properties of the application server in the content pane.
3. In the Additional Properties table, select **Message Listener Service**  
This displays the Message Listener Service properties in the content pane.
4. In the content pane, select **Listener Ports**.  
This displays a list of the listener ports.
5. **(Optional)** In the content pane, select the listener port that you want to stop.
6. Click **Stop**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

---

## Important files for message-driven beans and extended messaging

The following files in the WAS\_HOME/temp directory are important for the operation of the WebSphere Application Server messaging service, so should not be deleted. If you do need to delete the WAS\_HOME/temp directory or other files in it, ensure that you preserve the following files.

### *server\_name-durableSubscriptions.ser*

You should not delete this file, because the messaging service uses it to keep track of durable subscriptions for message-driven beans. If you uninstall an application that contains a message-driven bean, this file is used to unsubscribe the durable subscription.

### *server\_name-AsyncMessageRequestLog.ser*

You should not delete this file, because the messaging service uses it to keep track of late responses that need to be delivered to the late response message handler for the extended messaging provider.

---

## Troubleshooting message-driven beans

Use this overview task to help resolve a problem that you think is related to message-driven beans.

Message-driven beans support uses the standard WebSphere Application Server troubleshooting facilities. If you encounter a problem that you think might be related to the message-driven beans, complete the following stages:

### Steps for this task

1. Check for messages about message-driven beans in the application server's SystemOut log at `was_home\logs\server\SystemOut`.  
Look in the SystemOut log for messages that indicate a problem with JMS resources for message-driven beans, such as listener ports.
2. Check for more messages in the application server's SystemOut log.  
If the JMS server is running, but you have problems accessing JMS resources, check the SystemOut log file, which should contain more error messages and extra details about the problem.
3. Check the Release Notes for specific problems and workarounds  
The section *Possible Problems and Suggested Fixes* of the Release Notes, available from the WebSphere Application Server library web site, is updated regularly to contain information about known defects and their workarounds. Check the latest version of the Release Notes for any information about your problem. If the Release Notes does not contain any information about your problem, you can also search the Technotes database on the WebSphere Application Server web site.
4. Check that message listener service has started.  
The message listener service is an extension to the JMS functions of the JMS provider. It provides a listener manager that controls and monitors one or more JMS listeners, which each monitor a JMS destination on behalf of a deployed message-driven bean.
5. Check your JMS resource configurations  
If the WebSphere Messaging functions seem to be running properly (the JMS server is running without problems), check that the JMS resources have been configured correctly. For example, check that the listener ports have been configured correctly and have been started.
6. Check for problems with the WebSphere Messaging functions  
For more information about troubleshooting WebSphere Messaging, see the related topics.
7. **(Optional)** Get a detailed exception dump for messaging.  
If the information obtained in the preceding steps is still inconclusive, you can enable the application server debug trace for the "Messaging" group to provide a detailed exception dump.

---

## Message-driven beans samples

The following examples are provided to illustrate use of the message-driven beans support:

- Point-to-point samples:
  - "Tutorial: Creating JMS message sample"  
This tutorial is designed to help you develop and deploy a JMS message sample application that tests the WebSphere Application Server message-driven beans support in a point-to-point scenario. This sample illustrates how to develop and deploy an application that comprises the following components:
    - A Java/JMS program that writes a message to a queue.
    - A message-driven bean that is invoked by a JMS listener when a message arrives on a defined queue.

For more information about this sample, see the samples article "Tutorial: Creating JMS message sample" at `WAS_HOME/Enterprise/samples/messaging/doc/PtoP/wsbldPtoP.xml` (if you have installed the samples option).

- "Sample: Message Listener (point-to-point)"

This sample is designed to demonstrate the use and behavior of message-driven beans for a simple point-to-point scenario. This sample uses the JMS message sample deployed in the sample above.

For more information about this sample, see the samples article "Sample: Message Listener (Point-to-Point)" at `WAS_HOME/Enterprise/samples/messaging/doc/PtoP/wsrnPtoP.xml` (if you have installed the samples option).

- Publish/subscribe samples

- "Tutorial: Creating JMS message publish/subscribe sample"

This tutorial is designed to help you develop and deploy a JMS message sample application that tests the WebSphere Application Server message-driven beans support in a publish/subscribe scenario. This sample illustrates how to develop and deploy an application that comprises the following components:

- A client program that starts the message sequence by publishing a message to a selected topic.
- A message-driven bean that is invoked by a JMS listener when the broker passes a message to the listener from a topic to which it has subscribed.

For more information about this sample, see the samples article "Tutorial: Creating JMS message publish/subscribe sample" at `WAS_HOME/Enterprise/samples/messaging/doc/PubSub/wsbldPubSub.xml` (if you have installed the samples option).

- "Sample: Message Listener (publish/subscribe)"

This sample is designed to demonstrate the use and behavior of message-driven beans for a simple publish/subscribe scenario. This sample uses the JMS message sample deployed in the publish/subscribe sample above.

For more information about this sample, see the samples article "Sample: Message Listener (publish/subscribe)" at `WAS_HOME/Enterprise/samples/messaging/doc/PubSub/wsrnPUBSub.xml` (if you have installed the samples option).



---

## Chapter 7. Using application clients

An application client module is a JAR (Java ARchive) file containing a client for accessing a Java application.

### Steps for this task

1. Decide on a type of application client.
2. Develop the application client code.
  - Develop ActiveX application client code.
  - Develop applet client code.
  - Develop J2EE application client code.
  - Develop pluggable application client code.
  - Develop thin application client code.

### Usage scenario

View the Samples gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started.

These samples do not include client applications that communicate with Enterprise server-side examples.

---

## Application clients

In a traditional client server environment, the client requests a service and the server fulfills the request. Multiple clients use a single server. Clients can also access several different servers. This model persists for Java clients except now these requests make use of a client run-time environment.

In this model, the client application requires a servlet to communicate with the enterprise bean, and the servlet must reside on the same machine as the WebSphere Application Server.

With WebSphere Application Server V5.0, application clients now consist of the following models:

- ActiveX application client
- Applet client
- J2EE application client
- Pluggable application client
- Thin application client

The *ActiveX application client* model, uses the Java Native Interface (JNI) architecture to programmatically access the Java virtual machine (JVM) API. Therefore the JVM code exists in the same process space as the ActiveX application (Visual Basic, VBScript, or Active Server Pages (ASP)) and remains attached to the process until that process terminates.

In the *Applet client* model, a Java applet embeds in a HyperText Markup Language (HTML) document residing on a remote client machine from the WebSphere

Application Server. With this type of client, the user accesses an enterprise bean in the WebSphere Application Server through the Java applet in the HTML document.

The *J2EE application client* is a Java application program that accesses enterprise beans, Java Database Connectivity (JDBC), and Java Message Service message queues. The J2EE application client program runs on client machines. This program follows the same Java programming model as other Java programs; however, the J2EE application client depends on the application client run time to configure its execution environment, and uses the Java Naming and Directory Interface (JNDI) name space to access resources.

The *Pluggable and thin application clients* provide a lightweight Java client programming model. These clients are best suited in situations where a Java client application exists but the application needs enhancements to use enterprise beans, or where the client application requires a thinner, more lightweight environment than the one offered by the J2EE application client. The difference between the thin application client and the pluggable application client is that the thin application client includes a Java virtual machine (JVM) API, and the pluggable application client requires the user to provide this code. The pluggable application client uses the Sun Java Development Kit, and the thin application client uses the IBM Developer Kit For the Java Platform.

The J2EE application client programming model provides the benefits of the J2EE platform for the Java client application. The J2EE application client offers the ability to seamlessly develop, assemble, deploy and launch a client application. The tooling provided with the WebSphere platform supports the seamless integration of these stages to help the developer create a client application from start to finish.

When you develop a client application using and adhering to the J2EE platform, you can put the client application code from one J2EE platform implementation to another. The client application package can require redeployment using each J2EE platform deployment tool, but the code that comprises the client application does not change.

The application client run time supplies a container that provides access to system services for the client application code. The client application code must contain a main method. The application client run time invokes this main method after the environment initializes and runs until the Java virtual machine code terminates.

The J2EE platform allows the application client to use *nicknames* or *short names*, defined within the client application deployment descriptor. These deployment descriptors identify enterprise beans or local resources (JDBC, Java Message Service (JMS), JavaMail and URL APIs) for simplified resolution through JNDI use. This simplified resolution to the enterprise bean reference and local resource reference also eliminates changes to the client application code, when the underlying object or resource either changes or moves to a different server. When these changes occur, the application client can require redeployment.

The application client also provides initialization of the run-time environment for the client application. The deployment descriptor defines this unique initialization for each client application. The application client run time also provides support for security authentication to the enterprise beans and local resources.

The application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access enterprise bean references and to use CORBA services provided by the J2EE platform implementation. Use of the RMI-IIOP protocol and



the accessibility of CORBA services assist users in developing a client application that requires access to both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments or programming models in one client application, you must understand the differences between the two programming models to use and manage each appropriately.

View the Samples gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started.

These samples do not include client applications that communicate with Enterprise server-side examples.

## Application client functions

Use the following table to identify the available functions in the different types of clients.

| Available functions  | ActiveX client | Applet client | J2EE client | Pluggable client | Thin client |
|--|----------------|---------------|-------------|------------------|-------------|
| Provides all the benefits of a J2EE platform                                       | Yes            | No            | Yes         | No               | No          |
| Portable across all J2EE platforms   | No             | No            | Yes         | No               | No          |
| Provides the necessary run-time to support communication between client and server | Yes            | Yes           | Yes         | Yes              | Yes         |
| Allows the use of nicknames in the deployment descriptors                          | Yes            | No            | Yes         | No               | No          |
| Supports use of the RMI-IIOP protocol  | Yes            | Yes           | Yes         | Yes              | Yes         |
| Browser based application  | No             | Yes           | No          | No               | No          |

|  |     |         |     |     |     |
|--|-----|---------|-----|-----|-----|
| Enables development of client applications that can access enterprise bean references and CORBA object references            | Yes | Yes     | Yes | Yes | Yes |
| Enables the initialization of the client application run-time environment  | Yes | No      | Yes | No  | No  |
| Supports security authentication to enterprise beans   | Yes | Limited | Yes | Yes | Yes |
| Supports security authentication to local resources  | Yes | No      | Yes | No  | No  |
| Requires distribution of application to client machines  | Yes | No      | Yes | Yes | Yes |
| Enables access to enterprise beans and other Java classes through Visual Basic, VBScript, and Active Server Pages (ASP) code | Yes | No      | No  | No  | No  |
| Provides a lightweight client suitable for download  | No  | Yes     | No  | Yes | Yes |
| Enables access to Java Naming and Directory Interface (JNDI) for enterprise bean resolution                                  | Yes | Yes     | Yes | Yes | Yes |

|   |    |    |     |     |    |
|---|----|----|-----|-----|----|
| Runs on client machines that use the Sun Java Runtime Environment                                 | No | No | No  | Yes | No |
| Supports CORBA services (using CORBA services can render the application client code nonportable) | No | No | Yes | No  | No |

## ActiveX application clients

WebSphere Application Server provides an ActiveX to EJB bridge that enables ActiveX programs to access WebSphere Enterprise JavaBeans through a set of ActiveX automation objects.

The bridge accomplishes this by loading the Java virtual machine (JVM) into any ActiveX automation container such as Visual Basic, VBScript, and Active Server Pages (ASP).

There are two main environments in which the ActiveX to EJB bridge runs:

- **Client applications**, such as Visual Basic and VBScript, are programs that a user starts from the command line, desktop icon, or Start menu shortcut.
- **Client services**, such as Active Server Pages, are programs started by some automated means like the Services control panel applet.

The ActiveX to EJB bridge uses the Java Native Interface (JNI) architecture to programmatically access the JVM code. Therefore the JVM code exists in the same process space as the ActiveX application (Visual Basic, VBScript, or ASP) and remains attached to the process until that process terminates. To create JVM code, an ActiveX client program calls the `XJBInit()` method of the `XJB.JClassFactory` object. For more information about creating JVM code for an ActiveX program, see *ActiveX to EJB bridge, initializing JVM code*.

After an ActiveX client program has initialized the JVM code, the program calls several methods to create a proxy object for the Java class. When accessing a Java class or object, the real Java object exists in the JVM code; the automation container contains the proxy for that Java object. The ActiveX program can use the proxy object to access the Java class, object fields, and methods. For more information about using Java proxy objects, see *ActiveX to EJB bridge, using Java proxy objects*. For more information about calling methods and access fields, see *ActiveX to EJB bridge, calling Java methods and ActiveX to EJB bridge, accessing Java fields*.

The client program performs primitive data type conversion through the COM `IDispatch` interface (use of the `IUnknown` interface is not directly supported). Primitive data types are automatically converted between native Automation types and Java types. All other types are handled automatically by the Proxy Objects. For more information about data type conversion, see *ActiveX to EJB bridge, converting data types*.

Any exceptions thrown in Java code are encapsulated and re-thrown as a COM error, from which the ActiveX program can determine the actual Java exceptions. For more information about handling exceptions, see ActiveX to EJB bridge, handling errors.

The ActiveX to EJB bridge supports both free-threaded and apartment-threaded access and implements the Free Threaded Marshaler to work in a hybrid environment such as Active Server Pages. For more information about the support for threading, see ActiveX to EJB bridge, using threading.

## Applet clients

The applet client provides a browser-based Java run time capable of interacting with enterprise beans directly, instead of indirectly through a servlet.

This client is designed to support users who want a browser-based Java client application programming environment that provides a richer and more robust environment than the one offered by the Applet > Servlet > enterprise bean model.

The programming model for this client is a cross between the Java application thin client and a servlet client. When accessing enterprise beans from this client, the applet can consider the enterprise bean object references as CORBA object references.

No tooling support exists for this client to develop, assemble or deploy the applet. You are responsible for developing the applet, generating the necessary client bindings for the enterprise beans and CORBA objects, and bundling these pieces together to install or download to the client machine. The Java applet client provides the necessary run time to support communication between the client and the server.

Client side bindings generate using the Chapter 27, "Assembling applications", on page 1149. An applet can utilize these bindings, or you can generate client side bindings using the **rmic** command that is part of the IBM Developer Kit, Java edition, installed with the WebSphere Application Server.

The Applet client uses the RMI-IIOP protocol. Using this protocol enables the applet to access enterprise bean references and CORBA object references, but the applet is restricted in using some supported CORBA services.

If you combine the enterprise bean and CORBA environments in one applet, you must understand the differences between the two programming models, and you must use and manage each appropriately.

The applet client provides the run time to support the J2EE applet client. The applet client does not have tooling support for developing, assembling or deploying the applet. The applet client run time is provided through the Java applet browser plug-in that you install on the client machine using the WebSphere Application Server Client CD.

The applet environment restricts access to external resources from the browser run-time environment. You can make some of these resources available to the applet by setting the correct security policy settings in the WebSphere Application Server `client.policy` file. If given the correct set of permissions, the applet client must explicitly create the connection to the resource using the appropriate API.

This client does not perform initialization of any service that the client applet can need. For example, the client application is responsible for the initialization of the naming service, either through CosNaming or Java Naming and Directory Interface (JNDI) APIs.

## J2EE application clients

The J2EE application client programming model provides the benefits of Java™ 2 Platform for WebSphere Application Server Enterprise (J2EE).

The J2EE platform offers the ability to seamlessly develop, assemble, deploy and launch a client application. The tooling provided with the WebSphere platform supports the seamless integration of these stages to help the developer create a client application from start to finish.

When you develop a client application using and adhering to the J2EE platform, you can put the client application code from one J2EE platform implementation to another. The client application package can require redeployment using each J2EE platform deployment tool, but the code that comprises the client application does not change.

The J2EE application client run time supplies a container that provides access to system services for the application client code. The J2EE application client code must contain a main method. The J2EE application client run time invokes this main method after the environment initializes and runs until the Java virtual machine application terminates.

Application clients can use *nicknames* or *short names*, defined within the client application deployment descriptor with the J2EE platform. These deployment descriptors identify enterprise beans or local resources (Java Database Connectivity (JDBC), Java Message Service (JMS), JavaMail and URL APIs) for simplified resolution through JNDI use. This simplified resolution to the enterprise bean reference and local resource reference also eliminates changes to the application client code, when the underlying object or resource either changes or moves to a different server. When these changes occur, the application client can require redeployment.

The J2EE application client also provides initialization of the run-time environment for the client application. The deployment descriptor defines this unique initialization for each client application. The J2EE application client run time also provides support for security authentication to the enterprise beans and local resources.

The J2EE application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access enterprise bean references and to use CORBA services provided by the J2EE platform implementation. Use of the RMI-IIOP protocol and the accessibility of CORBA services assist users in developing a client application that requires access to both enterprise bean references and CORBA object references.

When you combine the J2EE and the CORBA WebSphere Application Server Enterprise environments or programming models in one client application, you must understand the differences between the two programming models to use and manage each appropriately.

## Pluggable application clients

The pluggable application client provides a lightweight, downloadable Java application run time capable of interacting with enterprise beans.

The pluggable application client requires that you have previously installed the Sun Java Runtime Environment (JRE) files. In all other aspects, the pluggable application client, and the thin application client are similar.

**Note:** *The pluggable client is only available on the Windows platform.*

This client is designed to support those users who want a lightweight Java client application programming environment, without the overhead of the J2EE platform on the client machine. The programming model for this client is heavily influenced by the CORBA programming model, but supports access to enterprise beans.

When accessing enterprise beans from this client, the client application can consider the enterprise beans object references as CORBA object references.

Tooling does not exist on the client; however, tooling does exist on the server. You are responsible for developing the client application, generating the necessary client bindings for the enterprise bean and CORBA objects, and after bundling these pieces together, installing them on the client machine.

The pluggable application client provides the necessary run time to support the communication needs between the client and the server.

The pluggable application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access not only enterprise bean references and CORBA object references, but the protocol also allows the client application to use any supported CORBA services. Using the RMI-IIOP protocol along with the accessibility of CORBA services can assist a user in developing a client application that needs to access both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments in one client application, you must understand the differences between the two programming models to use and manage each appropriately.

The pluggable application client run time provides the necessary support for the client application for object resolution, security, Reliability Availability and Serviceability (RAS), and other services. However, this client does not support a container that provides easy access to these services. For example, no support exists for using *nicknames* for enterprise beans or local resource resolution. When resolving to an enterprise bean (using either Java Naming and Directory Interface (JNDI) or CosNaming) sources, the client application must know the location of the name server and the fully qualified name used when the reference was bound into the name space. When resolving to a local resource, the client application cannot resolve to the resource through a JNDI lookup. Instead the client application must explicitly create the connection to the resource using the appropriate API (JDBC, Java Message Service (JMS), and so on). This client does not perform initialization of any of the services that the client application might require. For example, the client application is responsible for the initialization of the naming service, either through CosNaming or JNDI APIs.

The pluggable application client offers access to most of the available client services in the J2EE application client. However, you cannot access the services in

the pluggable client as easily as you can in the J2EE application client. The J2EE client has the advantage of performing a simple Java Naming and Directory Interface (JNDI) name space lookup to access the desired service or resource. The pluggable client must code explicitly for each resource in the client application. For example, looking up an enterprise bean Home requires the following code in a J2EE application client:

```
java.lang.Object ejbHome = initialContext.lookup("java:/comp/env/ejb/MyEJBHome");
MyEJBHome = (MyEJBHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, MyEJBHome.class);
```

However, you need more explicit code in a Java pluggable application client:

```
java.lang.Object ejbHome =
    initialContext.lookup("the/fully/qualified/path/to/actual/home/in/namespace/MyEJBHome");
MyEJBHome = (MyEJBHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, MyEJBHome.class);
```

In this example, the J2EE application client accesses a logical name from the `java:/comp` name space. The J2EE client run time resolves that name to the physical location and returns the reference to the client application. The pluggable client must know the fully qualified physical location of the enterprise bean Home in the name space. If this location changes, the pluggable client application must also change the value placed on the `lookup()` statement.

In the J2EE client, the client application is protected from these changes because it uses the logical name. A change can require a redeployment of the EAR file, but the actual client application code remains the same.

The pluggable application client is a traditional Java application that contains a *main* function. The WebSphere pluggable application client provides run time support for accessing remote enterprise beans, and provides the implementation for various services (security, Workload Management (WLM), and others). This client can also access CORBA objects and CORBA based services. When using both environments in one client application, you need to understand the differences between the enterprise bean and CORBA programming models to manage both environments.

For instance, the CORBA programming model requires the CORBA CosNaming name service for object resolution in a name space. The enterprise beans programming model requires the JNDI name service. The client application must initialize and properly manage these two naming services.

Another difference applies to the enterprise bean model. Use the Java Naming and Directory Interface (JNDI) implementation in the enterprise bean model to initialize the Object Request Broker (ORB). The client application is unaware that an ORB is present. The CORBA model, however, requires the client application to explicitly initialize the ORB through the `ORB.init()` static method.

The pluggable application client provides a batch command that you can use to set the `CLASSPATH` and `JAVA_HOME` environment variables to enable the pluggable application client run time.

## Thin application clients

The thin application client provides a lightweight, downloadable Java application run time capable of interacting with enterprise beans.

This client is designed to support those users who want a lightweight Java client application programming environment, without the overhead of the J2EE platform

on the client machine. The programming model for this client is heavily influenced by the CORBA programming model, but supports access to enterprise beans.

When accessing enterprise beans from this client, the client application can consider the enterprise beans object references as CORBA object references.

Tooling does not exist on the client, it exists on the server. You are responsible for developing the client application, generating the necessary client bindings for the enterprise bean and CORBA objects, and bundling these pieces together to install on the client machine.

The thin application client provides the necessary run-time to support the communication needs between the client and the server.

The thin application client uses the RMI-IIOP protocol. Using this protocol enables the client application to access not only enterprise bean references and CORBA object references, but also allows the client application to use any supported CORBA services. Using the RMI-IIOP protocol along with the accessibility of CORBA services can assist a user in developing a client application that needs to access both enterprise bean references and CORBA object references.

When you combine the J2EE and CORBA environments in one client application, you must understand the differences between the two programming models, to use and manage each appropriately.

The thin application client run time provides the necessary support for the client application for object resolution, security, Reliability Availability and Servicability (RAS), and other services. However, this client does not support a container that provides easy access to these services. For example, no support exists for using *nicknames* for enterprise beans or local resource resolution. When resolving to an enterprise bean (using either Java Naming and Directory Interface (JNDI) or CosNaming) sources, the client application must know the location of the name server and the fully qualified name used when the reference was bound into the name space. When resolving to a local resource, the client application cannot resolve to the resource through a JNDI lookup. Instead the client application must explicitly create the connection to the resource using the appropriate API (JDBC, Java Message Service (JMS), and so on). This client does not perform initialization of any of the services that the client application might require. For example, the client application is responsible for the initialization of the naming service, either through CosNaming or JNDI APIs.

The thin application client offers access to most of the available client services in the J2EE application client. However, you cannot access the services in the thin client as easily as you can in the J2EE application client. The J2EE client has the advantage of performing a simple Java Naming and Directory Interface (JNDI) name space lookup to access the desired service or resource. The thin client must code explicitly for each resource in the client application. For example, looking up an enterprise bean Home requires the following code in a J2EE application client:

```
java.lang.Object ejbHome = initialContext.lookup("java:/comp/env/ejb/MyEJBHome");
MyEJBHome = (MyEJBHome)javadb.rmi.PortableRemoteObject.narrow(ejbHome, MyEJBHome.class);
```

However, you need more explicit code in a Java thin application client:

```
java.lang.Object ejbHome =
    initialContext.lookup("the/fully/qualified/path/to/actual/home/in/namespace/MyEJBHome");
MyEJBHome = (MyEJBHome)javadb.rmi.PortableRemoteObject.narrow(ejbHome, MyEJBHome.class);
```



In this example, the J2EE application client accesses a logical name from the `java:/comp` name space. The J2EE client run time resolves that name to the physical location and returns the reference to the client application. The thin client must know the fully qualified physical location of the enterprise bean Home in the name space. If this location changes, the thin client application must also change the value placed on the `lookup()` statement.

In the J2EE client, the client application is protected from these changes because it uses the logical name. A change might require a redeployment of the EAR file, but the actual client application code remains the same.

The thin application client is a traditional Java application that contains a *main* function. The WebSphere thin application client provides run-time support for accessing remote enterprise beans, and provides the implementation for various services (security, Workload Management (WLM), and others). This client can also access CORBA objects and CORBA based services. When using both environments in one client application, you need to understand the differences between the enterprise bean and CORBA programming models to manage both environments.

For instance, the CORBA programming model requires the CORBA CosNaming name service for object resolution in a name space. The enterprise beans programming model requires the JNDI name service. The client application must initialize and properly manage these two naming services.

Another difference applies to the enterprise bean model. Use the Java Naming and Directory Interface (JNDI) implementation in the enterprise bean model to initialize the Object Request Broker (ORB). The client application is unaware that an ORB is present. The CORBA model, however, requires the client application to explicitly initialize the ORB through the `ORB.init()` static method.

The thin application client provides a batch command that you can use to set the `CLASSPATH` and `JAVA_HOME` environment variables to enable the thin application client run time.

---

## Example: Migrating application clients

Use the `ClientUpgrade` command to migrate application clients.

The command file is located in the `bin` subdirectory of the `<i>WAS_install_root</i>`, or the `<i>ND_install_root</i>` directory. By default, the `WAS_install_root` for WebSphere Application Server and WebSphere Application Server Enterprise is:

- Windows NT or Windows 2000 operating platforms - `<drive>\WebSphere\AppServer` directory
- AIX or UNIX-based operating platforms - `/usr/WebSphere/AppServer` directory

By default, the `ND_install_root` for WebSphere Application Server Network Deployment is:

- Windows NT or Windows 2000 operating platforms - `<drive>\WebSphere\DeploymentManager` directory
- AIX or UNIX-based operating platforms - `/usr/WebSphere/DeploymentManager` directory

The command uses the `com.ibm.websphere.migration.clientupgrade.ClientUpgrade` class.

## Syntax

```
ClientUpgrade EAR_file [-clientJar client_jar ]  
                    [-traceString trace_spec [-traceFile file_name ]]
```

## Parameters

Supported arguments include the following:

### EAR\_file

Use this parameter to specify the fully qualified path to the EAR file that contains client JAR files to process.

### -clientJar

Use this optional parameter to specify a JAR file for processing. If not specified, the program transforms all client JAR files in the EAR file.

### -traceString -traceFile

Use these optional parameters to gather trace information for IBM Service personnel. Specify a trace\_spec of "\*=all=enabled" (with quotation marks) to gather all trace information.

## Examples

The following example demonstrates correct syntax:

```
ClientUpgrade EAR_file -clientJar ejbJarFile
```

## Migration tips for application clients

Tips for migrating thin application client code:

- The Java invocation used to run non-J2EE application clients has changed in Version 5.0. You must specify `-Xbootclasspath/p:%WAS_BOOTCLASSPATH%` on Windows systems or `-Xbootclasspath/p:$WAS_BOOTCLASSPATH` on Unix systems when you invoke the Java command. Set the `WAS_BOOTCLASSPATH` environment variable in one of the following:
  - `setupClient.bat` for Windows systems or `setupClient.sh` for Unix systems on a WebSphere Application Server client installation.
  - `setupCmdLine.bat` for Windows systems or `setupCmdLine.sh` for Unix systems on a WebSphere Application Server installation.

For more information about using `-Xbootclasspath`, view sample code at the following path after you preform the application client installation:

```
install_root\samples\bin\ActiveXBridgeClients\VB\XJBExamples\modXJBHelpers.bas
```

Tips for migrating J2EE application client code:

- If your J2EE application client uses resource references and you have configured those resources using the Application Client Resource Configuration Tool (ACRCT), you must run the **ClientUpgrade** command to migrate the resource configuration information in WebSphere Application Server V5.

---

## Installing application clients

### Before you begin

Application clients do not require the WebSphere Application Server base. They are supported at the level of the operating system. If you want to run the Samples gallery, however, you must provide the host name of the server.

Before you install the pluggable application client, you must have one of the Sun Java Runtime Environment (JRE) on your system. Similar Sun Java Development Kit levels are required to install the pluggable application client Samples. *Although you can install the WebSphere Application Server on an HP\_UX system, the WebSphere application client is not supported on HP-UX systems.* See the Supported Prerequisites (<http://www.ibm.com/software/webservers/appserv/doc/latest/prereq.html>) page in the IBM external Web site for more information on supported product platforms.

**Note:** The installation panel displays that the amount of storage needed to install the J2EE/Thin application client is 40.4 megabytes. The actual amount of storage needed is 75 megabytes.

#### Steps for this task

1. Insert the WebSphere Application Server Client CD, according to platform.

2. Change to the following directory, according to platform:

aix for AIX platform

sun for Solaris Operating Environment

linuxi386 for Linux platform

nt for Windows platform

3. Issue the following command, according to platform, at a command line:

./install for AIX platform

./install for Solaris platform

./install for Linux platform

install.exe for Windows platform

The WebSphere Application Server client installation wizard appears.

4. If the default language is not correct, select a language from the drop-down menu and click **OK**.

If you have an earlier version of the WebSphere Application Server client installed on your machine, an information window appears.

5. Click **Yes to all** to overwrite this older version. A Welcome panel displays.

6. Click **Next**.

7. Read the license agreement, and select the radio button next to **I accept the terms in the license agreement**. Click **Next**.

8. Choose a setup type and click **Next**.

9. If you choose the Custom setup type, select the features you want to install and click **Next**.

If you choose the Applet client (Windows systems only) and you have an existing development kit or runtime environment installed, click **Yes** to overwrite the existing development kit or runtime environment registry entries. If you click **No**, you return to the Custom Features panel, where you can clear Applet client. You must use the Applet client with the runtime environment provided by WebSphere Application Server.

10. If you choose to install the Samples, click **Yes**, to install the Samples development environment.

11. Click **Next**, to install the WebSphere Application Server client to the default directory. Click **Browse...**, to install the WebSphere Application Server client to a different directory.

12. Click **Next**.

13. Enter the host name of the WebSphere Application Server machine.  
Specify a port number if you are not using the default port. Click **Next**. A verification panel displays.
14. Click **Next**, to install.  
A progress window appears.
15. Click **Finish** to exit the wizard, after the WebSphere Application Server client installs.

### Results

Application clients are installed on your machine.

## **Installing application clients on Version 9 of the Solaris Operating Environment**

### Before you begin

Follow this procedure when installing application clients with the embedded messaging client feature on Version 9 of the Solaris Operating Environment. You need not perform this procedure if you are not installing the embedded messaging client feature. If you are installing from the CD-ROM labelled **Application Clients 1.4**, you need not perform this procedure because the CD is current.

Solaris, Version 9 requires an update to the `wsmfuncs.SunOS` file before you install the application client from the Version 5 or Version 5.0.1 **Application Clients** CD-ROM. IBM provides the `mqpreinst.tar` file that has the `wsmfuncs.SunOS` file, a current `prereqChecker.xml` configuration file (that is not needed for client installation), and the `mq_preinstall.sh` script. The script copies the installation image from the product CD-ROM or a remote disk to a local directory where it applies the updated `wsmfuncs.SunOS` file. When the script finishes its work, you can begin the installation.

### **Manually updating the installation image**

You can also download the `wsmfuncs.SunOS` file if you prefer to copy the installation image and replace the file yourself.

When you copy the entire installation image from CD-ROM to a local directory, you copy the messaging directory from the product CD-ROM to a temporary directory, such as `/tmp/WebSphereTemp`, for example. Simply replace the `/tmp/WebSphereTemp/sun/messaging/wsmfuncs.SunOS` file.

After updating the installation image, you can begin the installation.

### **Determining if you have a problem**

If you do not complete this procedure but install the embedded messaging client feature on the Solaris 9 platform, the installation appears to complete without visible error but you cannot use any embedded messaging functionality.

### Steps for this task

1. Download the `mqpreinst.tar` file from the [ ].

Click **Support downloads** and limit the downloads search to Version 5.0.2 to locate the article about installing WebSphere Application Server products and clients on Version 9 of the Solaris Operating Environment. The article contains the mqpreinst.tar file.

2. Unpack the mqpreinst.tar file.  
Unpacking the file creates the /mqpreinstall directory.
3. Use the mq\_preinstall.sh script to copy and update the installation image.

The basic syntax of the script is:

```
mq_preinstall.sh -orig originalImageLocation -temp tempImageLocation
```

The -orig parameter identifies the root location of the original WebSphere Application Server client installation files.

The -temp parameter identifies the local target directory where the script copies the original installation image. The target directory should not exist. The script creates the target directory.

For example:

```
mq_preinstall.sh -orig /cdrom -temp /tmp/WebSphereTemp
```

The script copies the original installation image from the CD-ROM to the /tmp/WebSphereTemp directory and updates the backlevel /tmp/WebSphereTemp/sun/messaging/wsmfuncs.SunOS file.

4. Continue with the installation, by starting the installation wizard from the target directory.
5. View the mq\_prereq log in the /tmp directory to verify that the embedded messaging feature installed successfully.

The return code from the installation program indicates success. Installation of the embedded messaging feature is successful if the return code is zero. If the installation program exited with a non-zero return code, the installation of the embedded messaging feature is not successful. If unsuccessful, uninstall the client and complete the procedure described above to reinstall.

## Results

After installing the updated installation image, install Fix Pack 2.

## Usage scenario

You can invoke the mq\_preinstall.sh script with a relative or absolute file path, or from within the mqpreinst directory:

### **Absolute addressing example**

```
/tmp/mqpreinstall/mq_preinstall.sh -orig /cdrom -temp /tmp/WebSphereTemp
```

### **Relative addressing example**

```
./mqpreinstall/mq_preinstall.sh -orig /cdrom -temp /tmp/WebSphereTemp
```

### **Example from within the /mqpreinst directory**

```
cd mqpreinstall  
./mq_preinstall.sh -orig /cdrom -temp /tmp/WebSphereTemp
```

---

## Developing ActiveX application client code

### Before you begin

This topic provides an outline for developing an ActiveX program, such as Visual Basic, VBScript, and Active Server Pages, to use the WebSphere ActiveX to EJB bridge to access enterprise beans.

This topic assumes that you are familiar with ActiveX programming. You should also consider the information given in ActiveX to EJB bridge, good programming guidelines.

To use the ActiveX to EJB bridge to access a Java class, develop your ActiveX program to complete the following steps:

### Steps for this task

1. Create an instance of the XJB.JClassFactory object.
2. Create JVM code within the ActiveX program process, by calling the XJBInit() method of the XJB.JClassFactory object.  
After the ActiveX program has created an XJB.JClassFactory object and called the XJBInit() method, the JVM code is initialized and ready for use.
3. Create a proxy object for the Java class, by using the XJB.JClassFactory FindClass() and NewInstance() methods.  
The ActiveX program can use the proxy object to access the Java class, object fields, and methods.
4. Call methods on the Java class, using the Java method invocation syntax, and access Java fields as required.
5. Use the helper functions to do the conversion in cases where automatic conversion is not possible. You can convert between the following data types:
  - Java Byte and Visual Basic Byte
  - Visual Basic Currency types and Java 64-bit
6. Implement methods to handle any errors returned from the Java class. In Visual Basic or VBScript, use the Err.Number and Err.Description fields to determine the actual Java error.

### What to do next

After you develop the ActiveX client code, start the ActiveX application.

## Starting an ActiveX application

### Before you begin

To run an ActiveX client application that is to use the ActiveX to EJB bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run-time. This initial configuration sets up the environment within which the ActiveX client application can run.

To perform the required configuration, complete one or more of the following subtasks:

### Steps for this task

1. Starting an ActiveX application and configuring service programs

## 2. Starting an ActiveX application and configuring non-service programs

### Starting an ActiveX application and configuring service programs

#### Before you begin

To run an ActiveX service program such as Active Server Page (ASP) that is to use the ActiveX to EJB bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run-time. This configuration sets up the environment within which the ActiveX service program can run.

The XJB.JClassFactory must find the Java run-time Dynamic Link Library (DLL) when initializing. In a service program such as Internet Information Server you cannot specify a path for its processes independently; you must set the process paths in the system PATH variable. This limitation means that you can only have a single JVM version available on a machine using ASP.

To add the JRE directories to your System path, complete one of the following subtasks:

#### Steps for this task

1. On Windows 2000, complete the following substeps:
  - a. Open the Control Panel, then double-click the System icon.
  - b. Click the Advanced tab on the System Properties window.
  - c. Click **Environment Variables**.
  - d. Edit the **Path** variable in the System Variables window.
  - e. Add the following to the beginning of the path displayed in the Variable Value input box:

```
C:\WebSphere\AppClient\Java\jre\bin;C:\WebSphere\AppClient\Java\jre\bin\classic;
```

where C:\WebSphere\AppClient is the directory in which you installed the WebSphere Java client

- f. Click **OK** in the Edit System Variable window to apply the changes.
  - g. Click **OK** in the Environment Variables window.
  - h. Click **OK** in the System Properties window.
  - i. Restart Windows 2000.
2. On Windows NT, complete the following substeps:
  - a. Open the Control Panel, then double-click the System icon.
  - b. Click the Environment tab on the System Properties window.
  - c. In the System Variables window, edit the **Path** variable.
  - d. Add the following to the beginning of the path displayed in the Value input box:

```
C:\WebSphere\AppClient\Java\jre\bin;C:\WebSphere\AppClient\Java\jre\bin\classic;
```

Where C:\WebSphere\AppClient is the directory in which you installed the WebSphere Java client

- e. Click **Set** to apply the changes.
  - f. Click **OK**.
  - g. Restart Windows NT.

## What to do next

After you change the System PATH variable you must reboot the Internet Information Server machine so that Internet Information Server can see the change.

## **Starting an ActiveX application and configuring non-service programs**

### Before you begin

To run an ActiveX program initiated from an icon or command-line (a non-service program) that is to use the ActiveX to EJB bridge, you must perform some initial configuration to set appropriate environment variables and to enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java runtime. This uses a batch file to set up the environment within which the ActiveX program can run.

To perform the required configuration, complete the following steps:

### Steps for this task

1. **(Optional)** Edit the setupCmdLineXJB.bat file to specify appropriate values for the environment variables required by the ActiveX to EJB bridge. For more information about these environment variables, see ActiveX to EJB bridge, environment and configuration.

For more information about creating a JVM for an ActiveX program, see ActiveX to EJB bridge, initializing the JVM.

After the ActiveX program has created an XJB.JClassFactory object and called the XJBInit() method, the JVM is initialized and ready for use.

2. Start the ActiveX client application by using one of the following methods:
  - Use the launchClientXJB.bat file to start the application; for example:  
`launchClientXJB MyApplication.exe parm1 parm2`
  - or  
`launchClientXJB MyApplication.vbp`
  - Use the setupCmdLineXJB.bat file to create an environment in which the application can be run, then start the application from within that environment.

## **setupCmdLineXJB.bat, launchClientXJB.bat, and other ActiveX batch files**

This topic provides reference information about the aids that client applications and client services can use to access the ActiveX to EJB bridge. These enable the ActiveX to EJB bridge to find its XJB.JAR file and the Java run-time.

### **Location**

The include file is located in the `<i>was_client_home</i>\aspIncludes` directory. You can include the file into your ASP application with the following syntax in your ASP page:

```
<!-- #include virtual ="/WSASPIncludes/setupASPXJB.inc" -->
```

This assumes that you have created a virtual directory in Internet Information Server called **WSASPIncludes** that points to the `<i>was_client_home</i>\aspIncludes` directory.

### **Usage notes**



The following batch files are provided for client applications to use the ActiveX to EJB bridge:

- **setupCmdLineXJB.bat**

Sets the client environment variables.

- **launchClientXJB.bat**

Calls the setupCmdLineXJB.bat file and launches the application you specify as its arguments; for example:

```
launchClientXJB.bat myapp.exe parm1 parm2
```

or

```
launchClientXJB MyApplication.vbp
```

- **Active Server Pages (ASP) include file**

An include file is provided for ASP users to automatically set the following page-level (local) environment variables:

- **com\_ibm\_websphere\_javahome** Path to the Java run-time directory installed with the WebSphere Advanced Server Client.
- **com\_ibm\_websphere\_washome** Path to the WebSphere Advanced Server Client directory.
- **com\_ibm\_websphere\_namingfactory** Sets the Java java.naming.factory.initial system property.
- **com\_ibm\_websphere\_computername** (Optional) Name of the computer where the WebSphere Advanced Server Client is installed. If you intend to talk to a single specific computer, you are recommended to change this value to become the server name that you intend to access.

- **System Settings**

To enable the ActiveX to EJB bridge to access the Java run-time Dynamic Link Library (DLL), the following directories must exist in the system PATH environment variable:

```
was_client_home\java\jre\bin;was_client_home\java\jre\bin\classic
```

Where *was\_client\_home* is the name of the directory where you installed the WebSphere Application Server Client (for example, C:\WebSphere\AppClient).

**Note:** This technique enables only *one* Java run-time to activate on a machine, therefore all client services on that machine must use the same Java run-time. Client applications do not have this limitation because they each have their own private, non-system scope.

## JClassProxy and JObjectProxy classes

This topic provides reference information about the object classes of the ActiveX to EJB bridge.

JClassFactory is the object used to access the majority of JVM features. It handles JVM initialization, accessing classes and creating class instances (objects). The majority of tasks for accessing your Java classes and objects are handled with the JClassProxy and JObjectProxy objects.

- XJBInit(String astrJavaParameterArray())

Initializes the JVM environment using an array of strings that represent the command line parameters you would normally send to the java.exe file.

If you have invalid parameters in the XJBInit() string array, the following error results:

Error: 0x6002 "XJBNI::Init() Failed to create VM" when calling XJBInit()

If you have C++ logging enabled, the activity log displays the invalid parameter.

- `JClassProxy FindClass(String strClassName)`

Uses the current thread class loader to load the specified fully qualified class name and returns a `JClassProxy` object representing the Java Class object.

- `JObjectProxy NewInstance()`

Creates a Class instance for the specified `JClassProxy` object using the parameters supplied to call the Class Constructor. For more information about using `JMethodArgs`, see ActiveX to EJB bridge, calling Java methods.

`JObjectProxy`

```
NewInstance(JClassFactory obj, Variant vArg1, Variant vArg2, Variant vArg3,...)
JObjectProxy NewInstance(JClassFactory obj, JMethodArgs args)
```

- `JMethodArgs GetArgsContainer()`

Returns a `JMethodArgs` object (Class instance).

You can create a `JClassProxy` object from the `JClassFactory.FindClass()` method and also from any Java method call that would normally return a Java Class object. You can use this object as if you had direct access to the Java Class object. All of the class static methods and fields are accessible as are the `java.lang.Class` methods. In case of a clash between static method names of the reflected user class and those of the `java.lang.Class` (for example, `getName()`), the reflected static methods would execute first.

For example, the following is a static method called `getName()`. The `java.lang.Class` object also has a method called `getName()`:

– In Java:

```
class foo{
    foo(){};
    public static String getName(){return "abcdef";}
    public static String getName2(){return "ghijkl";}
    public String toString2(){return "xyz";}
}
```

– In Visual Basic:

```
...
Dim clsFoo as Object
set clsFoo = oXJB.FindClass("foo")
clsFoo.getName() ' Returns "abcdef" from the static foo class
clsFoo.getName2() ' Returns "ghijkl" from the static foo class
clsFoo.toString() ' Returns "class foo" from the java.lang.Class object.
oFoo = oXJB.NewInstance(clsFoo)
oFoo.toString() ' Returns some text from the java.lang.Object's
                ' toString() method which foo inherits from.
oFoo.toString2() ' Returns "xyz" from the foo class instance
```

You can create a `JObjectProxy` object from the `JClassFactory.NewInstance()` method, and can be created from any Java method call that would normally return a Class instance object. You can use this object as if you had direct access to the Java object and can access all the static methods and fields of the object. All of object instance methods and fields are accessible (including those accessible through inheritance).

The `JMethodArgs` object is created from the `JClassFactory.GetArgsContainer()` method. Use this object as a container for method and constructor arguments. You must use this object when overriding the object type when calling a method (for example, when sending a `java.lang.String JProxyObject` to a constructor that normally takes a `java.lang.Object` type).

There are two groups of methods to add arguments to the collection: **Add** and **Set**. You can use **Add** to add arguments in the order that they are declared. Alternatively, you can use **Set** to set an argument based on its position in the argument list (where the first argument is in position 1).

For example, if you had a Java Object Foo that took a constructor of Foo(int, String, Object), you could use a JMethodArgs object as shown in the following code extract:

```
...
Dim oArgs as Object
set oArgs = oXJB.GetArgsContainer()

oArgs.AddInt(CLng(12345))
oArgs.AddString("Apples")
oArgs.AddObject("java.lang.Object", oSomeJObjectProxy)

Dim clsFoo as Object
Dim oFoo as Object
set clsFoo = oXJB.FindClass("com.mypackage.foo")
set oFoo = oXJB.NewInstance(clsFoo, oArgs)

' To reuse the oArgs object, just clear it and use the add method
' again, or alternatively, use the Set method to reset the parameters
' Here, we will use Set
oArgs.SetInt(1, CLng(22222))
oArgs.SetString(2, "Bananas")
oArgs.SetObject(3, "java.lang.Object", oSomeOtherJObjectProxy)

Dim oFoo2 as Object
set oFoo2 = oXJB.NewInstance(clsFoo, oArgs)
```

- **AddObject (String strObjectName, Object oArg)**

Adds an arbitrary object to the argument container in the next available position, casting the object to the class name specified in the first parameter. Arrays are specified using the traditional [] syntax; for example:

```
AddObject("java.lang.Object[] []", oMy2DArrayOfFooObjects)
```

or

```
AddObject("int[]", oMyArrayOfInts)
```

- **AddByte (Byte byteArg)**

Adds a primitive byte value to the argument container in the next available position.

- **AddBoolean (Boolean bArg)**

Adds a primitive boolean value to the argument container in the next available position.

- **AddShort (Integer iArg)**

Adds a primitive short value to the argument container in the next available position.

- **AddInt (Long lArg)**

Adds a primitive int value to the argument container in the next available position.

- **AddLong (Currency cyArg)**

Adds a primitive long value to the argument container in the next available position.

- **AddFloat (Single fArg)**

Adds a primitive float value to the argument container in the next available position.

- **AddDouble (Double dArg)**  
Adds a primitive double value to the argument container in the next available position.
- **AddChar (String strArg)**  
Adds a primitive char value to the argument container in the next available position.
- **AddString (String strArg)**  
Adds the argument in string form to the argument container in the next available position.
- **SetObject (Integer iArgPosition, String strObjectTypeName, Object oArg)**  
Adds an arbitrary object to the argument container in the specified position casting it to the class name or primitive type name specified in the second parameter. Arrays are specified using the traditional [] syntax; for example:  
`SetObject(1, "java.lang.Object[] []", oMy2DArrayOfFooObjects)`  
  
or  
`SetObject(2, "int[]", MyArrayOfInts)`
- **SetByte (Integer iArgPosition, Byte byteArg)**  
Sets a primitive byte value to the argument container in the position specified.
- **SetBoolean (Integer iArgPosition, Boolean bArg)**  
Sets a primitive boolean value to the argument container in the position specified.
- **SetShort (Integer iArgPosition, Integer iArg)**  
Sets a primitive short value to the argument container in the position specified.
- **SetInt (Integer iArgPosition, Long lArg)**  
Sets a primitive int value to the argument container in the position specified.
- **SetLong (Integer iArgPosition, Currency cyArg)**  
Sets a primitive long value to the argument container in the position specified.
- **SetFloat (Integer iArgPosition, Single fArg)**  
Sets a primitive float value to the argument container in the position specified.
- **SetDouble (Integer iArgPosition, Double dArg)**  
Sets a primitive double value to the argument container in the position specified.
- **SetChar (Integer iArgPosition, String strArg)**  
Sets a primitive char value to the argument container in the position specified.
- **SetString (Integer iArgPosition, String strArg)**  
Sets a java.lang.String value to the argument container in the position specified.
- **Object Item(Integer iArgPosition)**  
Returns the value of an argument at a specific argument position.
- **Clear()**  
Removes all arguments from the container and resets the next available position to one.
- **Long Count()**  
Returns the number of arguments in the container.

## Java virtual machine initialization tips

Initialize the Java virtual machine (JVM) code with the ActiveX to enterprise JavaBeans bridge. For an ActiveX client program (Visual Basic, VBScript, or ASP) to access Java classes or objects, the first step that the program must do is to create JVM code within its process. To create JVM code, the ActiveX program calls the XJBInit() method of the XJB.JClassFactory object. When an XJB.JClassFactory object is created and the XJBInit() method called, the JVM is initialized and ready to use.

- To enable the XJB.JClassFactory to find the Java run-time Description Definition Language (DLL) when initializing, the JRE bin and bin\classic directories must exist in the system path environment variable.
- The XJBInit() method accepts only one parameter; an array of strings. Each string in the array represents a command line argument that for a Java program you would normally specify on the Java.exe command line. This string interface is used to set the classpath, stack size, heap size, and debug settings. You can get a listing of these parameters by typing `java -?` from the command line.
- If you set a parameter incorrectly, you receive a 0x6002 "Failed to initialize VM" error message.
- Due to the current limitations of Java Native Interface (JNI), you cannot unload or reinitialize JVM code after it has loaded. Therefore, after XJBInit() has been called once, subsequent calls have no effect other than to create a duplicate JClassFactory object for you to access. It is best to store your XJB.JClassFactory object globally and continue to reuse that object.
- The following Visual Basic extract shows an example of initializing JVM code:

```
Dim oXJB as Object
set oXJB = CreateObject("XJB.JClassFactory")
Dim astrJavaInitProps(0) as String
astrJavaInitProps(0) = _
    "-Djava.class.path=.;c:\myjavaclasses;c:\myjars\myjar.jar"
oXJB.XJBInit(astrJavaInitProps)
```

## Example: Developing ActiveX to enterprise bean bridge, using Java proxy objects

To use Java proxy objects with the ActiveX to enterprise JavaBeans bridge:

- After an ActiveX client program (Visual Basic, VBScript, or ASP) has initialized the XJB.JClassFactory (and thereby the JVM), it can access Java classes and initialize Java objects. To do this, the client program uses the XJB.JClassFactory FindClass() and NewInstance() methods.
- In Java programming there are two ways to access Java classes: direct invocation through the Java compiler, and through the Java Reflection interface. Because the ActiveX to Java bridge needs no compilation and is a complete runtime interface to Java, it depends on the latter Reflection interface to access its classes, objects, methods, and fields. The XJB.JClassFactory FindClass() and NewInstance() methods behave very similarly to the Java Class.forName() and the Method.invoke() and Field.invoke() methods.
- XJB.JClassFactory.FindClass() takes the fully-qualified class name as its only parameter and returns a Proxy Object (JClassProxy). You can use the returned Proxy object like a normal Java Class object and call static methods and access static fields. You can also create a Class Instance (or object) from it, as described below. For example, the following Visual Basic code extract returns a Proxy object for the Java class `java.lang.Integer`:

```
...
Dim clsMyString as Object
Set clsMyString = oXJB.FindClass("java.lang.Integer")
```

- After the proxy is created, you can access its static information directly. For example, you can use the following code extract to convert a decimal integer to its hexadecimal representation.

```
...
Dim strHexValue as String
strHexValue = clsMyString.toHexString(CLng(255))
```

- The equivalent Java syntax is: `static String toHexString(int i)`. Because ints in Java programming are really 32-bits (which translates to Long in VB), the `CLng()` function converts the value from the default int to a long. Also, even though the `toHexString()` function returns a `java.lang.String`, the code extract does not return an Object Proxy. Instead, the returned `java.lang.String` is automatically converted to a native Visual Basic String.

To create an object from a class, you use the `JClassFactory.NewInstance()` method. This method creates an Object Instance and takes whatever parameters your Class Constructor needs. Once the object is created, you have access to all of its public instance methods and fields. For example, you can use the following Visual Basic code extract to create an instance of `java.lang.Integer`:

```
...
Dim oMyInteger as Object
set oMyInteger = oXJB.NewInstance(CLng(255))

Dim strMyInteger as String
strMyInteger = oMyInteger.toString
```

## Example: Calling Java methods in the ActiveX to enterprise bean bridge

In the ActiveX to EJB bridge, methods are called using the native language method invocation syntax.

- The following are important differences between Java invocation and ActiveX Automation invocation:
  - Unlike Java methods, ActiveX does not support method (and constructor) polymorphism; that is, you cannot have two methods in the same class with the same name.
  - Java methods are case-sensitive, but ActiveX Automation is not case-sensitive.
- You should take care when invoking Java methods through ActiveX Automation. If you use the wrong case on a method call or use the wrong parameter type, you get an Automation Error 438 "Object doesn't support this property or method" thrown.
- To compensate for Java polymorphic behavior, give the exact parameter types to the method call. The parameter types determine the correct method to invoke. For a listing of correct types to use, see ActiveX to EJB bridge, converting data types.
- For example, the following Visual Basic code would fail if `CLng()` was not present or `toHexString` was incorrectly typed as `ToHexString`:

```
...
Dim strHexValue as String
strHexValue = clsMyString.toHexString(CLng(255))
```

- Sometimes it is difficult to force some development environments to leave the case of your method calls unchanged. For example, in Visual Basic if you want to call a method `close()` (uncapitalized), Visual Basic would try to capitalize it "Close()". In Visual Basic, the only way to effectively get around this behavior is to use the `CallByName()` method. For example:

```
o.Close(123) 'Incorrect...
CallByName(o, "close", vbMethod, 123) 'Correct...
```

or in VBScript, use the Eval function:

```
o.Close(123)                'Incorrect...
Eval("o.Close(123)")       'Correct...
```

- The return value of a function is always converted dynamically to the correct type. However, you must take care to use the **set** keyword in Visual Basic. If you expect a non-primitive data type to return, you must use **set**. (If you expect a primitive data type to return, you do not need to use **set**.) For example:

```
Set oMyObject = o.GetObject
iMyInt = o.GetInt
```

- In some cases, you might not know the type of object returning from a method call, because wrapper classes are converted automatically to primitives (for example, java.lang.Integer returns an ActiveX Automation Long). In such cases, you might need to use your language built-in exception handling techniques to try to coerce the returned type (for example, On Error and Err.Number in Visual Basic).

- Methods with Character Arguments

Because ActiveX automation does not natively support character types supported by Java methods, the ActiveX to EJB bridge uses strings (byte or VT\_I1 do not work, because characters have multiple bytes in Java). If you try to call a method that takes a char or java.lang.Character type you must use the JMethodArgs argument container to pass character values to methods or constructors. For more information about how this argument container is used, see Methods with "Object" Type as Argument and Abstract Arguments.

- Methods with "Object" Type as Argument and Abstract Arguments

Because of the polymorphic nature of Java programming, the ActiveX to Java bridge uses direct argument type mapping to find a method. This works well in most cases, but sometimes methods are declared with a Parent or Abstract Class as an argument type (for example, java.lang.Object). You need the ability to send an object of arbitrary type to a method. To acquire this ability, you must use the XJB.JMethodArgs object to coerce your parameters to match the parameters on your method. You can get a JMethodArgs instance by using the JClassFactory.GetArgsContainer() method.

The JMethodArgs object is a container for method parameters or arguments. This container enables you to add parameters to it one-by-one and then you can send the JMethodArgs object to your method call. The JClassProxy and JObjectProxy objects recognize the JMethodArgs object and attempt to find the correct method and let the Java language coerce your parameters appropriately.

For example, to add an element to a Hashtable object the method syntax is Object.put(Object key, Object value). In Visual Basic, the method usage looks like this:

```
Dim oMyHashtable as Object
Set oMyHashtable = _
    oXJB.NewInstance(oXJB.FindClass("java.util.Hashtable"))
```

```
' This line will not work. The ActiveX to EJB bridge cannot find a method
' called "put" that has a short and String as a parameter:
```

```
oMyHashtable.put 100, "Dogs"
oMyHashtable.put 200, "Cats"
```

```
' You must use a XJB.JMethodArgs object instead:
```

```
Dim oMyHashtableArgs as Object
Set oMyHashtableArgs = oXJB.GetArgsContainer
oMyHashtableArgs.AddObject("java.lang.Object", 100)
oMyHashtableArgs.AddObject("java.lang.Object", "Dogs")
```

```
oMyHashtable.put oMyHashtableArgs
```

```
' Reuse the same JMethodArgs object by clearing it.
oMyHashtableArgs.Clear
oMyHashtableArgs.AddObject("java.lang.Object", 200)
oMyHashtableArgs.AddObject("java.lang.Object", "Cats")

oMyHashtable.put oMyHashTableArgs
```

## Java field programming tips

Using the ActiveX to EJB bridge to access Java fields has the same case sensitivity issue that it has when invoking methods. Field names must use the same case as the Java field syntax.

- Visual Basic has the same problem with unsolicited case changing on fields as it does with methods. For more information about this problem, see “Example: Calling Java methods in the ActiveX to enterprise bean bridge” on page 264. You might need to use the CallByName() function to set a field in the same way that you would call a method in some cases. For Fields, you use VBLet for primitive types and VBSet for Objects. For example:

```
o.MyField = 123 'Incorrect...
CallByName(o, "MyField", vbLet, 123) 'Correct...
```

or in VBScript:

```
o.MyField = 123 'Incorrect...
Eval("o.myField = 123") 'Correct...
```

## ActiveX to Java primitive data type conversion values

All primitive Java data types are automatically converted to native ActiveX Automation types. However, not all Automation data types are converted to Java types (for example, VT\_DATE). Variant data types are used for data conversion. Variant data types are a requirement of any Automation interface, and are used automatically by Visual Basic and VBScript. The tables below provide details about how primitive data types are converted between Automation types and Java types.

### ActiveX to Java primitive data type conversion

| Visual Basic Type | Variant Type | Java Type        | Notes  |
|-------------------|--------------|------------------|--|
| Byte              | VT_I1        | byte             | Byte in Visual Basic is unsigned, but is signed in Java data type. |
| Boolean           | VT_BOOL      | boolean          | None   |
| Integer           | VT_I2        | short            | None   |
| Long              | VT_I4        | int              | None   |
| Currency          | VT_CY        | long             | None   |
| Single            | VT_R4        | float            | None   |
| Double            | VT_R8        | double           | None   |
| String            | VT_BSTR      | java.lang.String | None   |
| String            | VT_BSTR      | char             | None   |
| Date              | VT_DATE      | N/A              | None   |

### Example: Using helper methods for data type conversion

Generally, data type conversion between ActiveX (Visual Basic and VBScript) and Java methods occurs automatically, as described in ActiveX to EJB bridge,



converting data types. However, the following helper functions are provided for cases where automatic conversion is not possible:

- Byte helper function
- Currency helper function
- Byte helper function

Because the Java Byte data type is signed (-127 through 128) and the Visual Basic Byte data type is unsigned (0 through 255), convert unsigned Bytes to a Visual Basic Integer, which look like the Java signed byte. To make this conversion, you can use the following helper function:

```
Private Function GetIntFromJavaByte(Byte jByte) As Integer
    GetIntFromJavaByte = (CInt(jByte) + 128) Mod 256 - 128
End Function
```

- Currency helper function

Visual Basic 6.0 cannot properly handle 64-bit integers like Java methods can (as the Long data type). Therefore, Visual Basic uses the Currency type, which is intrinsically a 64-bit data type. The only side effect of using the Currency type (the Variant type VT\_CY) is that a decimal point is inserted into the type. To extract and manipulate the 64-bit Long value in Visual Basic, use code like the following example. For more details on this technique for converting Currency data types, see Q189862, "HOWTO: Do 64-bit Arithmetic in VBA", on the Microsoft Knowledge Base.

```
' Currency Helper Types
Private Type MungeCurr
    Value As Currency
End Type
Private Type Munge2Long
    LoValue As Long
    HiValue As Long
End Type

' Currency Helper Functions
Private Function CurrToText(ByVal Value As Currency) As String
    Dim Temp As String, L As Long
    Temp = Format$(Value, "#.0000")
    L = Len(Temp)
    Temp = Left$(Temp, L - 5) & Right$(Temp, 4)
    Do While Len(Temp) > 1 And Left$(Temp, 1) = "0"
        Temp = Mid$(Temp, 2)
    Loop
    Do While Len(Temp) > 2 And Left$(Temp, 2) = "-0"
        Temp = "-" & Mid$(Temp, 3)
    Loop
    CurrToText = Temp
End Function

Private Function TextToCurr(ByVal Value As String) As Currency
    Dim L As Long, Negative As Boolean
    Value = Trim$(Value)
    If Left$(Value, 1) = "-" Then
        Negative = True
        Value = Mid$(Value, 2)
    End If
    L = Len(Value)
    If L < 4 Then
        TextToCurr = CCur(IIf(Negative, "-0.", "0.") & _
            Right$("0000" & Value, 4))
    Else
        TextToCurr = CCur(IIf(Negative, "-", "") & _
            Left$(Value, L - 4) & "." & Right$(Value, 4))
    End If
End Function
```

```

' Java Long as Currency Usage Example
Dim LC As MungeCurr
Dim L2 As Munge2Long

' Assign a Currency Value (really a Java Long)
' to the MungeCurr type variable
LC.Value = cyTestIn

' Coerce the value to the Munge2Long type variable
LSet L2 = LC

' Perform some operation on the value, now that we
' have it available in two 32-bit chunks
L2.LoValue = L2.LoValue + 1

' Coerce the Munge value back into a currency value
LSet LC = L2
cyTestIn = LC.Value

```

## Array tips for ActiveX application clients

Arrays are very similar between Java and Automation Containers like Visual Basic and VBScript. Here are some important points to consider when passing arrays back and forth between these containers:

- Java arrays cannot mix types. All Java arrays contain a single type, so when passing arrays of Variants to a Java Array, you must make sure that all of the elements in the Variant array are of the same base type. For example, in Visual Basic:

```

...
Dim VariantArray(1) as Variant
VariantArray(0) = CLng(123)
VariantArray(1) = CDb1(123.4)
oMyJavaObject.foo(VariantArray) ' Illegal!

VariantArray(0) = CLng(123)
VariantArray(1) = CLng(1234)
oMyJavaObject.foo(VariantArray) ' This works

```

- Arrays of Primitive Types are converted using the rules defined in Primitive Data Type Conversion.
- Arrays of Java Objects are handled through arrays of JObjectProxy objects.
- Arrays of JObjectProxy objects must be fully-initialized and of the correct associated Java type. When initializing an array in Visual Basic (for example, Dim oJavaObjects(1) as Object), you must set each object to a JObjectProxy before you send the array to Java. The bridge is unable to determine the type of null or empty Object values.
- When receiving an array from a Java method, the lower-bound is always zero. Java methods only support zero-based arrays.
- Nested or multi-dimensional arrays are treated as zero-based multi-dimensional arrays in Visual Basic and VBScript.
- Uninitialized arrays or Array Types are unsupported. When calling a Java method that takes an array of objects as a parameter, you must fully initialize the array of JObjectProxy objects.

## Error handling codes for ActiveX application clients

All exceptions thrown in Java code are encapsulated and re-thrown as a COM error through the ISupportErrorInfo interface and the EXCEPINFO structure of IDispatch::Invoke(); the Err object in Visual Basic and VBScript. Because there are

no error numbers associated with Java exceptions, whenever a Java exception is thrown, the entire stack trace is stored in the error description text and the error number assigned is 0x6003.

In Visual Basic or VBScript, you need to use the Err.Number and Err.Description fields to determine the actual Java error. Non-Java errors are thrown as you would expect via the IDispatch interface; for example, if a method cannot be found, then error 438 "Object doesn't support this property or method" is thrown.

| Error number | Description  |
|--------------|--|
| 0x6001       | Java Native Interface (JNI) error                          |
| 0x6002       | Initialization error                                       |
| 0x6003       | Java exception. Error description is the Java Stack Trace. |
| 0x6FFF       | General Internal Failure                                   |

## Threading tips

The ActiveX to EJB bridge supports both free-threaded and apartment-threaded access and implements the Free Threaded Marshaler to work in a hybrid environment such as Active Server Pages. Each thread created in the ActiveX process is mirrored in the Java environment when the thread communicates through the ActiveX to EJB bridge. In addition, once all references to Java objects (there are no JObjectProxy or JClassProxy objects) are loaded in an ActiveX thread, the ActiveX to EJB bridge detaches the thread from the JVM code. Therefore, you must be careful that any Java code that you access from a multi-threaded Windows application is thread-safe. Visual Basic and VBScript applications are both essentially single-threaded. Therefore, Visual Basic and VBScript applications do not have threading issues in the Java programs they access. Active Server Pages and multi-threaded C and C++ programs can have issues.

Consider the following scenario:

1. A multi-threaded Windows Automation Container (our ActiveX Process) starts. It exists on Thread A.
2. The ActiveX Process initializes the ActiveX to EJB bridge, which starts the JVM code. The JVM attaches to the same thread and internally calls it Thread1.
3. The ActiveX Process starts two threads: B and C.
4. Thread B in the ActiveX Process uses the ActiveX to EJB bridge to access an object that was created in Thread A. The JVM attaches to thread B and calls it Thread 2.
5. Thread C in the ActiveX Process never talks to the JVM code, so the VM never needs to attach to it. This is a case where the JVM code does not have a one-to-one relationship between ActiveX threads and Java threads.
6. Thread B later releases all of the JObjectProxy and JClassProxy objects that it used. The Java Thread 2 is detached.
7. Thread B again uses the ActiveX to EJB bridge to access an object that was created in Thread A. The JVM code attaches again to the thread and calls it Thread 3.

| ActiveX Process         | JVM Access by ActiveX Process |
|-------------------------|-------------------------------|
| Thread A - Created in 1 | Thread 1 - Attached in 2      |

| ActiveX Process         | JVM Access by ActiveX Process                                    |
|-------------------------|--|
| Thread B - Created in 4 | Thread 2 - Attached in 4, detached in 6 Thread 3 - Attached in 7 |
| Thread C - Created in 4 | Not available  |

## Threads and Active Server Pages

Active Server Pages (ASP) in Microsoft's Internet Information Server is a multi-threaded environment. When you create the XJB.JClassFactory object, you can store it in the Application collection as an Application-global object. All threads within your ASP environment can now access the same ActiveX to EJB bridge object. Active Server Pages by default creates 10 Apartment Threads per ASP process per CPU. This means that when your ActiveX to EJB bridge object is initialized any of the 10 threads can call this object, not just the thread that created it.

If you need to simulate single-apartment behavior, you can create a Single-Apartment Threaded ActiveX Dynamic Link Library (DLL) in Visual Basic and encapsulate the ActiveX to EJB bridge object there. This encapsulation guarantees that all access to the JVM object is on the same thread. You need to use the <OBJECT> tag to assign the XJB.JClassFactory to an Application object and must be aware of the consequences of introducing single-threaded behavior to a Web application.

The Microsoft KnowledgeBase has several articles about ASP and threads, including:

- Q243543 INFO: Do Not Store STA Objects in Session or Application
- Q243544 INFO: Component Threading Model Summary Under Active Server Pages
- Q243548 INFO: Design Guidelines for VB Components Under ASP

## Example: Viewing System.out message

The ActiveX to EJB bridge does not have a console available to view Java System.out messages. To view these messages when running a stand-alone client program (such as Visual Basic), you need to redirect the output to a file. For example:

```
launchClientXJB.bat MyProgram.exe > output.txt
```

- To view the System.out messages when running a Service program such as Active Server Pages, you need to override the Java System.out OutputStream object to FileOutputStream. For example, in VBScript:

```
'Redirect system.out to a file
' Assume that oXJB is an initialized XJB.JClassFactory object
Dim clsSystem
Dim oOS
Dim oPS
Dim oArgs

' Get the System class
Set clsSystem = oXJB.FindClass("java.lang.System")

' Create a FileOutputStream object
' Create a PrintStream object and assign to it our FileOutputStream
Set oArgs = oXJB.GetArgsContainer oArgs.AddObject "java.io.OutputStream", oOS
```

```

Set oPS = oXJB.NewInstance(oXJB.FindClass("java.io.PrintStream"), oArgs)

' Set our System OutputStream to our file
clsSystem.setOut oPS

```

## Example: Enabling logging and tracing for application clients

The ActiveX to EJB bridge provides two logging and tracing formats: Windows Application Event Log and Java Trace Log.

- Windows Application Event Log

The Windows Application Event Log shows JNI errors, Java console error messages, and XJB initialization messages. This log is most useful for determining `XJBInit()` errors and any unusual exceptions that do not come from the Java environment. By default, critical error logging will be enabled and debug and event logging will be disabled.

To enable or disable logging certain types of events to the Windows Application Event Log, you need to specify one or more parameters to `XJBInit()`. If more than one parameter is set, they will be processed in the order in which they appear in the input string array to `XJBInit()`. Once `XJBInit()` is initialized, these parameters can no longer be set/reset for the life of the process. Using Java `java.lang.System.setProperty()` to set these values also will have no effect.

- `-Dcom.ibm.ws.client.xjb.native.logging.debug=enabled|disabled`

Enables or disables debug level messages from appearing in the Windows operating system event log. This level of logging is most useful and shows most internal errors, user programming issues or configuration problems.

- `-Dcom.ibm.ws.client.xjb.native.logging.event=enabled|disabled`

Enables or disables event level messages from appearing in the Windows operating system event log.

- `-Dcom.ibm.ws.client.xjb.native.logging.*=enabled|disabled`

Enables or disables both event and debug level messages from appearing in the Windows operating system event log. It is not possible to disable some critical error messages from being displayed in the error log. Only debug and event level messages can be disabled.

Viewing the Windows application event log with the event viewer:

To open the event viewer in the Windows operating system, click **Start > Settings > Control Panel**. Double-click **Administrative Tools**, and then double-click **Event Viewer**. All ActiveX to EJB bridge events will have the text "WebSphere XJB" in the source column and will appear in the Application log. For information about using Event Viewer, click the **Action** menu in Event Viewer, and then click **Help**.

To open the even viewer in the Windows operating system, click **Start > Programs > Administrative Tools > Event Viewer**. All ActiveX to EJB bridge events have the text "WebSphere XJB" in the source column and display in the Application log. For information about using Event Viewer, click the **Help** menu in Event Viewer.

- Java Trace Log

The Java trace log displays information that you can use to debug method calls, class lookups, and argument coercion problems. Since the Java portion of the bridge mirrors the functionality of the COM IDispatch interface, the information

in the trace log is similar to what you have come to expect from an IDispatch interface. To understand the trace log, you need a fundamental understanding of IDispatch.

To enable user-logging, add the following parameters to the XJBInit() input string array:

```
"-DtraceString=com.ibm.ws.client.xjb.*=event=enabled"  
"-DtraceFile=C:\MyTrace.txt"
```

## ActiveX client programming best practices

In general, the best way to access Java components is to use the Java language. It is recommended that you do as much programming as possible in the Java language and use a small simple interface between your COM Automation container (for example, Visual Basic) and the Java code. This interface avoids any overhead and performance problems that can occur when moving across the interface.

- Visual Basic guidelines
- Active Server Pages guidelines
- J2EE guidelines

### Visual Basic guidelines

The following guidelines are intended to help optimize your use of the ActiveX to EJB bridge with Visual Basic:

- Launch the Visual Basic replication through the `launchClientXJB.bat` file. If you want to run your Visual Basic application through the Visual Basic debugger, run the Visual Basic Integrated Development Environment (IDE) within the ActiveX to EJB bridge environment. After you create your Visual Basic project, you can launch it from a command line; for example, `launchClientXJB MyApplication.vbp`. You can also launch the Visual Basic application alone in the ActiveX to EJB environment, by changing the Visual Basic shortcut on the Windows Start menu so that the `launchClientXJB.bat` file precedes the call to the `VB6.EXE` file.
- Exit the Visual Basic IDE before debugging programs.  
Because the Java virtual machine (JVM) code attaches to the running process, you must exit out of the Visual Basic editor before debugging your program. If you run then exit your program within the Visual Basic IDE, the JVM code continues to run and you reattach the same JVM code when `XJBInit()` is called by the debugger. This causes problems if you try to update `XJBInit()` arguments (for example, `classpath`) because the changes are not be applied until you restart Visual Basic.
- Store the `XJB.JClassFactory` object globally.  
Because you cannot unload or reinitialize the JVM code, cache the resulting `XJB.JClassFactory` object as a global variable. The overhead of treating this object as a global variable or passing a single reference around is much less than recreating a new `XJB.JClassFactory` object and calling the `XJBInit()` argument more than once.

### CScript and Windows Scripting Host

The following guidelines intend to help optimize your use of the ActiveX to EJB bridge with CScript and Windows Scripting Host (WSH):

- Launch in ActiveX to EJB environment.  
Launch the VBScriptfiles in the ActiveX to EJB bridge environment, to run VBScript files in `.vbs` files. Two common ways exist to launch your script:

- launchClientXJB MyScript.vbs
- launchClientXJB cscript MyScript.vbs

### Active Server Pages guidelines

The following guidelines intend to help optimize your use of the ActiveX to EJB bridge with Active Server Pages software:

- Use the ActiveX to EJB Helper functions from the Active Server Pages Application.

Because Active Server Pages (ASP) code typically use VBScript, you can use the included helper functions in any VBScript environment with minor changes. For more information about these helper functions, see *Helper functions for data type conversion*. To run outside of the ASP environment, remove or change all references to the Server, Request, Response, Application and Session objects; for example, change `Server.CreateObject` to `CreateObject`.

- Set JRE path globally in system.

The XJB.JClassFactory object must be able to find the Java run time Dynamic Link Library (DLL) when initializing. In Internet Information Server, you cannot specify a path for its processes independently; you must set the process paths in the system PATH variable. You can only have a single JVM version available on a machine using the ASP application. Also, remember that after you change the system PATH variable you must reboot the Internet Information Server machine so that Internet Information Server can see the change.

- Set the system TEMP environment variable.

If the system TEMP environment variable is not set, Internet Information Server stores all temporary files in the WINNT directory, which is usually not desired.

- Use high isolation or an isolated process.

When using the ActiveX to Java bridge with Active Server Pages software, creating your Web application in its own process is recommended. You can only load one JVM instruction in a single process and if you want to have more than one application running with different JVM environment options (for example, different classpaths), then you need to have separate processes.

- Use the Application Unload option.

When debugging your application, use **Unload** when viewing your ASP application properties in the Internet Information Server administration console to unload the process from memory and thereby unload the JVM code.

- Run one process per application.

Use only one ASP application per J2EE application or JVM environment, in your ASP environment. If you need separate classpaths or JVM settings you need separate ASP applications (virtual directories with high isolation or an isolated process).

- Store the XJB.JClassFactory object in application scope.

Because of the one-to-one relationship required between a JVM instruction and a process, and because the JVM code can never detach or shut down from a process independently, cache the XJB.JClassFactory object at application scope and call the XJBInit() method only once.

Because the ActiveX to EJB bridge employs a free-threaded marshaler, take advantage of the multi-threaded nature of Internet Information Server and the ASP environment. If you choose to reinitialize the XJB.JClassFactory object at Page scope (local variables), then the XJBInit() method can only initialize your local XJB.JClassFactory variable. It is more efficient to use the XJBInit() method once.

- Use VBScript conversion functions.  
Because VBScript code only supports variant data types, use the CStr(), CByte(), CBool(), CCur(), CInt(), CInq(), CSng() and CDb1() functions to tell the activeX to EJB bridge which data type you are using; for example `oMyObject.Foo(CDb1(1.234))`.

## J2EE guidelines

The following guidelines are intended to help optimize your use of the ActiveX to EJB bridge with the J2EE environment;

- Store client container objects globally.

Because you can only have one JVM instruction per process, and a single J2EE client container (`com.ibm.websphere.client.applicationclient.launchClient`) per JVM instruction, initialize your J2EE client container only once and reuse it. For ASP applications, store the J2EE client container in an application level variable and initialize it only once (either on the `Application_OnStart()` event in the `global.asa` file or by checking to see if it `IsEmpty()`).

A side effect to storing the client container object globally is that you cannot change the client container parameters without destroying the object and creating a new one. These parameters include the EAR file, `BootstrapHost`, `classpath`, and so on. If you run a Visual Basic application and want to change the client container parameters, you must end the application and restart it. If you run an Active Server Pages application, you must first unload the application from Internet Information Server (see "Use the Application Unload Button" under Active Server Pages guidelines). Then load the Active Server Pages application with the different client container parameters. The parameters set the first time the Active Server Pages application loads. Since the client container is stored on the Internet Information Server, all the browser clients share the parameters using the Active Server Pages application. This behavior is normal for Active Server Pages code, but can be confusing when you try to run to different WebSphere Application Servers using the same Active Server Pages applicatio, which is unsupported.

- Reuse custom temp directory for EAR file extraction.

By default, the client container launches and extracts the application EAR file to your temp directory and then sets up the thread `ClassLoader` to use the extracted EAR file directory and JAR files included in the client JAR manifest. This process is time consuming and because of some limitations with JVM shutdown through Java Native Interface (JNI) and file locking, these files are never cleaned up.

Specifically, each time the client container `launch()` method is called, it extracts the EAR file to a random directory name in your temporary directory on your hard drive. The current Java thread class loader is then changed to point to this extracted directory which in turn locks the files within. In a normal J2EE Java client, these files automatically clean up after the application exits. This cleanup occurs when the client container shutdown hook is called (which never happens in the ActiveX to EJB bridge), which leaves the temporary directory there.

To avoid these problems, you can specify a directory to extract the EAR file by setting the `com.ibm.websphere.client.applicationclient.archivedir` Java system property before calling the client container `launch()` method. If the directory does not exist or is empty, you extract the EAR file normally. If the EAR file was previously extracted, the directory is reused. This feature is particularly important for server processes (for example, ASP), which can stop and restart, potentially calling `launchClient()` several times.



If you need to update your EAR file, delete the temporary directory first. The next time you create the client container object, it extracts the new EAR file to the temporary directory. If you do not delete the temporary directory or change the system property value to point to a different temporary directory, the client container reuses the currently extracted EAR file, and does not use your changed EAR file.

**Note:** When specifying the `com.ibm.websphere.client.applicationclient.archivedir` property, make sure that the directory you specify *is unique* for each EAR file you use. For example, do not point `MyEar1.ear` and `MyEar2.ear` files to the same directory.

If you choose not to use this system property, go regularly to your Windows temp directory and delete the WSTMP\* subdirectories. Over a relatively short period of time, these subdirectories can waste a very significant amount of space on the hard drive.

---

## Developing applet client code

### Before you begin

Applet clients have the following setup requirements:

- These clients are currently available on the Windows NT or Windows 2000 platforms. Check the prerequisites page for information on new platform support.
- They require one of these browsers:
  - Internet Explorer version 5.0+
  - Netscape Navigator 4.7+
- You must install the browser before installing the client code.

Unlike typical applets that reside on either Web servers or WebSphere Application Servers and can only communicate using the HTTP protocol, applet clients are capable of communicating over the HTTP protocol and the RMI-IIOP protocol. This additional capability gives the applet direct access to enterprise beans.

### Steps for this task

1. Run the application server client installation.
2. Select the applet client option.
3. Install an applet client.
4. Install the WebSphere Application Server Plug-in for the browser.

From the WebSphere Application Server Java Plug-in Control panel, enter the following:

```
-Djava.security.policy=product_installation_dir\properties\client.policy
-Dwas.install.root=product_installation_dir
-Djava.ext.dirs=product_installation_dir\classes;product_installation_dir
\java\jre\lib\ext;product_installation_dir\java\jre\lib;product_installation_dir
\lib;product_installation_dir\properties
-Dcom.ibm.CORBA.securityEnabled=false
-Dcom.ibm.CORBA.ConfigURL=file:product_installation_dir\properties\sas.client.props
-classpath product_installation_dir\properties
```

**Note:** The above entries are automatically placed into the WebSphere Application Server Java Plug-in control panel for the user who installed the WebSphere Application Server client. If this sample is being run by a user other than the person who installed the client, the user must enter the entries.

- The *Java Run Time Parameters* field is similar to the command prompt when using command line options. Therefore, you can enter most options available from the command prompt (for example, `-cp`, `classpath`, and others) in this field as well.
- Access the control panel from the **Start** menu. Click **start > Control panel > WebSphere Java Plug-in**.
- The applet container is the Web browser and the Java plug-in combination. You must first install the WebSphere Application Server Applet client so that the browser recognizes the IBM Java Plug-in.

View the Samples gallery for more information about application clients. Before you run the `basicCalculator` Sample, ensure the JMS Server is started.

These samples do not include client applications that communicate with Enterprise server-side examples.

## Accessing secure resources using the TCP/IP protocol for applet clients

You cannot use the WebSphere Application Server JSSE from an applet, and applets cannot authenticate to WebSphere Application Server using Secure Sockets Layer (SSL). To access resources that are not secure, such as enterprise beans, applets are configured with security disabled. Applets can access secure resources using standard TCP/IP protocol. This will cause the userid and password to flow from the client machine to the application server in the clear. This may compromise your password should someone be monitoring the network. It is important to understand the security implications before you modify the client configuration to use the standard TCP/IP protocol to access secure resources on your application server.

### Steps for this task

1. Make a copy of the following file so that you can use it for an applet:

```
<product_install_directory>/properties/sas.client.props
```

It is recommended to use a copy of the `sas.client.props` file for your applet because this file is used for all WebSphere Application Server client applications.

2. Edit the copy of `sas.client.props` file that you made with the following changes:
  - Change `com.ibm.CSI.protocol` to `com.ibm.CSI.protocol=csiv2`.
  - Change `com.ibm.CSI.performTransportAssocSSLTLSSupported` to `com.ibm.CSI.performTransportAssocSSLTLSSupported=false`.
  - Add `com.ibm.CSI.claimTransportAssocSSLTLSSupported=false`.
3. Click **Start > Control panel > WebSphere Java Plug-in** to open the WebSphere Application Server Java control panel.
  - Change `-Dcom.ibm.CORBA.securityEnabled=false` to `-Dcom.ibm.CORBA.securityEnabled=true`.
  - To use the file you created in step number 1, modify the value:
 

```
-Dcom.ibm.CORBA.ConfigURL=file:<product_install_directory>\properties\sas.client.props
```

For more information on the `sas.client.props` file and WebSphere Application Server security, view [Security](#).

## Applet client security requirements

When code is loaded, it is assigned permissions based on the security policy in effect. This policy specifies the permissions that are available for code from various locations. You can initialize this policy from an external policy file. By default, the client uses the `<product_installation_dir>/properties/client.policy` file. You must update this file with the following permissions:

- The `SocketPermission` grants permission to open a port and make a connection to a host machine, which is your WebSphere Application Server. In the following example, `yourserver.yourcompany.com` is the complete hostname of your WebSphere Application Server:

```
permission java.util.PropertyPermission "*", "read";
permission java.net.SocketPermission "yourserver.yourcompany.com", "connect";
```

## Applet client tag requirements

Standard applets require the HTML `<APPLET>` tag to identify the applet to the browser. The `<APPLET>` tag invokes the Java Virtual Machine (JVM) of the browser.

- For applets to communicate with EJBs in the WebSphere Application Server environment, the `<APPLET>` tag must be replaced with the following tags:

```
<OBJECT>
<EMBED>
```

- The `classid` and `type` attributes cannot be modified, and must be entered as described in the applet client example. The `codebase` attribute on the `<OBJECT>` tag must be excluded. Do not confuse the `codebase` attribute on the `<OBJECT>` tag with the `codebase` attribute on the `<PARAM>` tag. Although both are called `codebase`, they are separate entities.
- The following code example illustrates the applet code. In this example, `MyApplet.class` is the applet code, `applet.jar` is the file that contains the applet code, and `EJB.jar` is the file that contains the enterprise bean code:

```
<OBJECT classid="clsid:8AE2D840-EC04-11D4-AC77-006094334AA9"
width="600" height="500">
<PARAM NAME=CODE VALUE=MyAppletClass.class>
<PARAM NAME="archive" VALUE='Applet.jar, EJB.jar'>
<PARAM TYPE="application/x-java-applet;version=1.3">
<PARAM NAME="scriptable" VALUE="false">
<PARAM NAME="cache-option" VALUE="Plugin">
<PARAM NAME="cache-archive" VALUE="Applet.jar, EJB.jar">
<COMMENT>
<EMBED type="application/x-websphere-client" CODE=MyAppletClass.class
ARCHIVE="Applet.jar, EJB.jar" WIDTH="600" HEIGHT="500"
scriptable="false">
<NOEMBED>
</COMMENT>
</NOEMBED>WebSphere Java Application/Applet Thin Client for
Windows is required.
</EMBED>
</OBJECT>
```

- The value of the `type` attribute on the `</EMBED>` tag can also be, for example:  
`<EMBED type="application/x-websphere-client, version=4.0" ...`

## Applet client code requirements

The code used by an applet to talk to an enterprise bean is the same as that used by a standalone Java program or a servlet, except for one additional property called `java.naming.applet`. This property informs the `InitialContext` and the `Object Request Broker (ORB)` that this client is an applet rather than a standalone Java application or servlet.

- When you initialize an instance of the InitialContext class, the first two lines in this code snippet illustrate what both a standalone Java program and a servlet issue to specify the computer name, domain, and port. In this example, <yourserver.yourdomain.com> is the computer name and domain where WebSphere Application Server resides, and 900 is the configured port. After the bootstrap values (<yourserver.yourdomain.com>:900) are defined, the client to server communications occur within the underlying infrastructure. In addition to the first two lines, for applets, you must add the highlighted third line to your code. That line identifies this program as an applet, for example:
 

```
prop.put(Context.INITIAL_CONTEXT_FACTORY,
"com.ibm.websphere.naming.WsnInitialContextFactory");
prop.put(Context.PROVIDER_URL, "iiop://<yourserver.yourdomain.com>:900)
prop.put(Context.APPLET, this);
```

---

## Developing J2EE application client code

A *J2EE application client* program operates similarly to a standard J2EE program in that it runs its own Java Virtual Machine code and is invoked at its main method.

The Java Virtual Machine application client program differs from a standard Java program because it uses the Java Naming and Directory Interface (JNDI) name space to access resources. In a standard Java program, the resource information is coded in the program.

Storing the resource information separately from the client application program makes the client application program portable and more flexible.

### Steps for this task

#### 1. Writing the client application program

Write the J2EE application client program on any development machine. At this stage, you do not require access to the WebSphere Application Server.

Using the `javax.naming.InitialContext` class, the client application program uses the lookup operation to access the Java Naming and Directory Interface (JNDI) name space. The `InitialContext` class provides the lookup method to locate resources.

The following example illustrates how a client application program uses the `InitialContext` class:

```
import javax.naming.*

public class myAppClient
{
    public static void main(String argv[])
    {
        InitialContext initCtx = new InitialContext();
        Object homeObject = initCtx.lookup("java:comp/env/ejb/BasicCalculator");
        BasicCalculatorHome bcHome =
            (BasicCalculatorHome) javax.rmi.PortableRemoteObject
.narrow(homeObject, BasicCalculatorHome.class);
        BasicCalculatorHome bc = bcHome.create();
        ...
    }
}
```

In this example, the program looks up an enterprise bean called `BasicCalculator`. The `BasicCalculator` EJB reference is located in the client JNDI name space at `java:comp/env/ejb/BasicCalculator`. Since the actual

enterprise bean runs on the server, the application client run time returns a reference to the BasicCalculator home interface.

If the client application program lookup was for a resource reference or an environment entry, then lookup returns an instance of the configured type as defined by the client application deployment descriptor. For example, if the program lookup was a JDBC datasource, the lookup would return an instance of `javax.sql.DataSource`.

## 2. Assemble the application client using the Application Assembly Tool

The JNDI name space knows what to return on a lookup because of the information assembled by the Application Assembly Tool (AAT).

Assemble the J2EE application client on any development machine with the AAT installed.

When you use the Application Assembly Tool to assemble your application client, you provide the *application client* run time with the required information to initialize the execution environment for your client application program. Refer to the Application Assembly Tool description for implementation details.

Following is a list of things to keep in mind when you configure resources used by your client application program:

- When configuring resource references, resource environment references, and EJB references in the Application Assembly Tool, the General tab contains a required Name field. This field specifies where the application client run time binds the reference to the real object in the `java:comp/env` portion of the JNDI name space. The application client run time always binds these references relative to `java:comp/env`. For the programming example above, specify `ejb/BasicCalculator` in the Name field on the General tab of the Application Assembly Tool, which requires the program to perform a lookup of `java:comp/env/ejb/BasicCalculator`. If the Name field is set to `myString`, the resulting lookup is `java:comp/env/myString`.
- When configuring Resource references in the Application Assembly Tool, the Name field on the General tab is used for:
  - Binding a reference of that object type into the JNDI name space.
  - Retrieving client specific resource configuration information that was configured using the Application Client Resource Configuration Tool.
- When configuring a resource reference in the Application Assembly Tool, the value in the Name field on the General tab must match the value in the JNDI Name field on the General tab for the resource in the Application Client Resource Configuration Tool.
- When configuring URL resources using the Application Client Resource Configuration Tool, the URL provider panel enables you to specify a protocol and a class to handle that protocol. If you want to use the default protocols, such as HTTP, you can leave those fields blank.
- When configuring resource references using the Application Assembly Tool, the General tab contains a field called Authorization. You can set this field to either Container or Application. If you set the field to Container, then the application client run time uses authorization information configured in the Application Client Resource Configuration tool for the resource. If the field is set to Application, then the application client run time expects the user application to provide authorization information for the resource. The application client run time ignores any authorization information configured with the Application Client Resource Configuration tool for that resource.
- When configuring resource environment references using the Application Assembly Tool, you must specify the location of the actual object in the

server JNDI namespace using the Binding tab. A resource environment reference maps a logical name (the Name field on the general tab) used by the client application to the physical name of an object (the JNDI Name field on the Bindings tab). Not all objects bound into the server JNDI namespace are intended for use by an application client. For example, the WebSphere Application Server client run-time does not support the use of Java 2 Connector (J2C) objects on the client. The object needs to be remotable, and the client-side implementations must be made available on the application client run-time classpath.

- Resource environment references are different than resource references. Resource environment references allow your application client to use a logical name to look-up a resource bound into the server JNDI namespace. A resource reference allows your application to use a logical name to look-up a local J2EE resource. The J2EE specification does not specify a particular implementation of a resource. The following is a table of the supported resource types and identifies the resources to which the WebSphere Application Server provides a client implementation.

| Resource Type  | Client Configuration Notes  | Client implementation provided by WebSphere Application Server |
|--|---|--|
| javax.sql.DataSource   | Supports specification of any Datasource implementation class                                 | No   |
| java.net.URL   | Supports specification of custom protocol handlers  | Provided by Java Runtime Environment files                     |
| javax.mail.Session   | Supports custom protocol configuration  | Yes - POP3, SMTP, IMAP   |
| javax.jms.QueueConnectionFactory, javax.jms.TopicConnectionFactory, javax.jms.Queue, javax.jms.Topic | Supports configuration of WebSphere Embedded Messaging, IBM MQ Series and other JMS providers | Yes - WebSphere Embedded Messaging                             |

### 3. Assembling the Enterprise Archive (EAR)

The application is contained in an enterprise archive or .ear file. The .ear file is composed of:

- Enterprise bean, application client, and user-defined modules or .jar files
- Web applications or .war files
- Metadata describing the applications or application .xml files

You must assemble the .ear file on the server machine.

### 4. Distribute the EAR file

The client machines configured to run this client must have access to the .ear file.

If all the machines in your environment share the same image and platform, run the Application Client Resource Configuration Tool (ACRCT) on one machine to configure the external resources, and then distribute the configured .ear file to the other machines.

If your environment is set up with a variety of client installations and platforms, run the ACRCT for each unique configuration.

You can either distribute the .ear files to the correct client machines, or make them available on a network drive.

Distributing the .ear files is the responsibility of the system and network administrator.

5. Deploy the application client.
6. Configure the application client resources

If the client application defines the local resources, run the ACRCT (clientConfig command) on the local machine to reconfigure the .ear file. Use the ACRCT to change the configuration. For example, the .ear file can contain a DB2 resource, configured as C:\DB2. If, however, you installed DB2 in the D:\Program Files\DB2 directory, use the ACRCT to create a local version of the .ear file.

### What to do next

After developing the J2EE application client code, launch the application client.

## **J2EE application client class loading**

When you run your J2EE application client, a hierarchy of class loaders is created to load classes used by your application.

The following list describes the hierarchy of class loaders:

- The topmost class loader, the bootstrap class loader, contains the JAR files that make up the Java Virtual Machine code, such as rt.jar, plus those JAR files defined by the -Xbootclasspath parameter on the Java command. The WebSphere Application client run time sets this value to the WAS\_BOOTCLASSPATH environment variable.
- The *extensions class loader* class loader is a child to the bootstrap class loader. This class loader contains JAR files in the java/jre/lib/ext directory or those JAR files defined by the -Djava.ext.dirs parameter on the Java command. The WebSphere Application Client run time does not set -Djava.ext.dirs parameters, so it uses the JAR files in the java/jre/lib/ext directory.
- The *system class loader* class loader contains JAR files and classes that are defined by the -classpath parameter on the java command. The Application Client run time sets this parameter to the WAS\_CLASSPATH environment variable.
- The *WebSphere class loader* class loader loads the WebSphere Application Client run time and any classes placed in the WebSphere Application Client user directories. The directories used by this class loader are defined by the WAS\_EXT\_DIRS environment variable. The WAS\_BOOTCLASSPATH, WAS\_CLASSPATH, and the WAS\_EXT\_DIRS environment variables are set in the installation\_root/bin/setupCmdLine command shell for WebSphere Application Server server installations, or in the installation\_root/bin/setupClient command shell for client installations.

As the J2EE application client run time initializes, additional class loaders are created as children of the WebSphere class loader. If your client application uses resources such as Java Database Connectivity (JDBC) API, Java Message Service (JMS) API, or Uniform Resource Locator (URL), a different class loader is created to load each of those resources. Finally, the application client run time sets the WebSphere class loader to load classes within the .ear file by processing the client JAR manifest repeatedly. The system classpath, defined by the CLASSPATH environment variable is never used and is not part of the hierarchy of class loaders.

To package your client application correctly, you must understand which class loader loads your classes. When Java loads a class, the class loader used to load

that class is assigned to it. Any classes subsequently loaded by that class will use that class loader or any of its parents, but it will not use children class loaders.

In some cases the WebSphere Application Client run time can detect when your client application class is loaded by a different class loader from the one created for it by the WebSphere Application Client run time. When this detection occurs, you see the following message:

```
WSCL0205W: The incorrect class loader was used to load [0]
```

This message occurs when your client application class is loaded by one of the parent class loaders in the hierarchy. This situation is typically caused by having the same classes in the .ear file and on the hard drive. If one of the parent class loaders locates a class, that class loader loads it before the application client run time class loader. In some cases, your client application will still function correctly. In most cases, however, you receive "class not found" exceptions.

### Configuring the classpath fields

When packaging your J2EE client application, you must configure various classpath fields. Ideally, you should package everything required by your application into your .ear file. This is the easiest way to distribute your J2EE client application to your clients. However, you should not package such resources as JDBC APIs, JMS APIs, or URLs. In the case of these resources, use classpath references to access those classes on the hard drive. You might also have other classes installed on your client machines that you do not need to redistribute. In this case, you also want to use classpath references to access the classes on the hard drive, as described below.

### Referencing classes within the EAR file

WebSphere J2EE applications do not use the system class path. Use the MANIFEST Classpath entry to refer to other JARs within the .ear file. Configure these values using the module Classpath fields in the Application Assembly Tool. For example, if your client application needs to access the path of the enterprise bean JAR, add the deployed enterprise bean module name to your application client Classpath field in the Application Assembly Tool. The format of the Classpath field for each of the different modules (Application Client, enterprise bean, Web) is the same:

- The values must refer to .jar and .class files that are contained within the .ear file.
- The values must be relative to the root of the .ear file.
- The values cannot refer to absolute paths in the file systems.
- Multiple values must be separated by spaces, not colons or semi-colons.

**Note:** This is the Java method for allowing applications to function platform-independent.

Typically, you add modules (.jar files) to the root of the .ear file. In this case, you only need to specify the name of the module (.jar file) in the Classpath field. If you choose to add a module with a path, you need to specify the path relative to the root of the .ear file.

For referencing .class files, you must specify the directory relative to the root of the .ear file. With the Application Assembly Tool you can add individual class files to the .ear file. It is recommended that these additional class files are packaged in a .jar file. Add this .jar file to the module Classpath fields. If you



add .class files to the root of the .ear file, add ./ to the module Classpath fields. Consider the following example directory structure in which the file myapp.ear contains an application client JAR file named client.jar and a mybeans.jar EJB module. Additional classes reside in class1.jar and utility/class2.zip files. A class named xyz.class is not packaged in a JAR file but is in the root of the EAR file. Specify ./ **mybeans.jar utility/class2.zip class1.jar** as the value of the Classpath property. The search order is: myapp.ear/client.jar myapp.ear/xyz.class myapp.ear/mybeans.jar myapp.ear/utility/class2.zip myapp.ear/class1.jar

### Referencing classes that are not in the EAR file

Use the launchClient -CCclasspath parameter. This parameter is specified at run time and takes platform-specific classpath values, which means multiple values are separated by semi-colons or colons. There are many similarities between the client and the server in this respect.

### Resource classpaths

When you configure resources used by your client application using the Application Client Resource Configuration Tool, you can specify classpaths that are required by the resource. For example, if your application is using a JDBC to a DB2 database, add db2java.zip to the classpath field of the database provider. These classpath values are platform-specific and require semi-colons or colons to separate multiple values.

### Using the launchClient API

If you use the **launchClient** shell and bat command, the WebSphere class loader hierarchy is created for you. However, if you use the launchClient API, you must perform this setup yourself. You should mimic the launchClient shell command in defining the Java system properties.

---

## Developing pluggable application client code

### Before you begin

As you prepare to install the pluggable application client, remember that pluggable clients are only available on Windows systems.

### Steps for this task

1. Install the pluggable application client from the WebSphere Application Client CD by selecting option **Pluggable Application Client** from the **Custom client installation** panel.
2. Set the Java application pluggable client environment by using the **setupClient** shell, located in:  
install\_root\AppClient\bin\setupClient.bat
3. Add your specific Java client application JAR files to the CLASSPATH and start your Java client application from this environment, after setting the environment variables.
4. Run the following Java command to invoke your client application:  
%JAVA\_HOME%/bin/java -Xbootclasspath/p:%WAS\_BOOTCLASSPATH%  
-classpath *list\_of\_your\_application\_jars\_and\_classes*  
-Djava.ext.dirs= %WAS\_EXT\_DIRS%

```

-Djava.naming.provider.url=iiop://your_WebSphere_server_machine_name
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
%SERVER_ROOT% %CLIENTSAS% fully_qualified_class_name_to_run
$JAVA_HOME/bin/java -Xbootclasspath/p:$WAS_BOOTCLASSPATH
-classpath list_of_your_application_jars_and_classes
-Djava.ext.dirs=$WAS_EXT_DIRS
-Djava.naming.provider.url=iiop://your_WebSphere_server_machine_name
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
$SERVER_ROOT $CLIENTSAS fully_qualified_class_name_to_run

```

### Usage scenario

View the Samples gallery for more information about application clients. Before you run the basicCalculator Sample, ensure the JMS Server is started.

These samples do not include client applications that communicate with Enterprise server-side examples.

---

## Developing thin application client code

You can develop and run Java thin client applications on machines installed with either a client or a server. The client provides a setup command shell which sets up your environment for either a thin client application or a J2EE client application. The server provides a command shell which sets up your environment for J2EE application clients only. The Java invocation to run a thin application client varies between a client and a server. If your thin client application needs to run on both a client installation and a server installation, follow the steps for developing thin application clients on a server machine.

### Steps for this task

1. Install the Java application thin client from the WebSphere Application Client CD by selecting option **J2EE/Thin application client for the WebSphere Application Client**.
2. Perform one of the following:
  - Develop thin application client code for a client machine.
  - Develop thin application client code for a server machine.

### Usage scenario

View the Samples gallery for more information about application clients.

These samples do not include client applications that communicate with Enterprise server-side examples.

## Developing thin application client code on a client machine

### Before you begin

You must install the thin application client from the WebSphere Application Client CD before performing this task. For more information, see Developing thin application client code.

### Steps for this task

1. Set the Java application thin client environment by using the **setupClient** shell, located in:

```
install_root\AppClient\bin\setupClient.bat (on Windows)
install_root/AppClient/bin/setupClient.sh (on UNIX platforms)
```

2. Run the following Java command to invoke your client application:

On Windows:

```
"%JAVA_HOME%\bin\java" "-Xbootclasspath/p:%WAS_BOOTCLASSPATH%"
-Djava.security.auth.login.config="%WAS_HOME%\properties\wsjaas_client.conf"
-classpath "%WAS_CLASSPATH%;<list of your application jars and classes>"
-Djava.ext.dirs=%WAS_EXT_DIRS% -Djava.naming.provider.url=<an iiop URL
or a corbaloc URL to your Websphere server machine name>
-Djava.naming.factory.initial=
com.ibm.websphere.naming.WsnInitialContextFactory
"%SERVER_ROOT%" "%CLIENTSAS%" <fully qualified class name to run><your application parameters>
```

On Unix:

```
$JAVA_HOME/bin/java -Xbootclasspath/p:$WAS_BOOTCLASSPATH
-classpath "$WAS_CLASSPATH:<list of your application jars and classes>"
-Djava.ext.dirs=$WAS_EXT_DIRS -Djava.naming.provider.url=<an iiop URL
or a corbaloc URL to your Websphere server machine name>
-Djava.naming.factory.initial=
com.ibm.websphere.naming.WsnInitialContextFactory
"$SERVER_ROOT" "$CLIENTSAS" <fully qualified class name to run><your application parameters>
```

For more information on iiop and corbaloc URLs, see “Developing applications that use JNDI” on page 1056.

### Usage scenario

View the Samples gallery for more information about application clients.

These samples do not include client applications that communicate with Enterprise server-side examples.

## Developing thin application client code on a server machine

### Before you begin

You must install the thin application client from the WebSphere Application Client CD before performing this task. For more information, see Developing thin application client code.

### Steps for this task

1. Set the Java application thin client environment by using the **setupCmdLine** shell, located in:

```
install_root\bin\setupCmdLine.bat (on Windows)
install_root/bin/setupCmdLine.sh (on UNIX platforms)
```

2. Run the application client. Perform one of the following methods:

- Use Java to call your main class directly:

On Windows:

```
"%JAVA_HOME%\bin\java" "-Xbootclasspath/p:%WAS_BOOTCLASSPATH%"
-Djava.security.auth.login.config="%WAS_HOME%\properties\wsjaas_client.conf"
-Djava.ext.dirs="%JAVA_HOME%\jre\lib\ext;%WAS_EXT_DIRS%"
-Djava.naming.provider.url=<an iiop URL or a corbaloc URL to your
Websphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
-Dserver.root="%WAS_HOME%" "%CLIENTSAS%" %USER_INSTALL_PROP%
-classpath "%WAS_CLASSPATH%;<list of your application jars and classes>"
<fully qualified class name to run><your application parameters>
```

On Unix:

```
"$JAVA_HOME/bin/java" "-Xbootclasspath/p:$WAS_BOOTCLASSPATH"
-Djava.security.auth.login.config="$WAS_HOME/properties/wsjaas_client.conf"
-Djava.ext.dirs="$JAVA_HOME/jre/lib/ext;%WAS_EXT_DIRS%"
-Djava.naming.provider.url=<an iiop URL or a corbaloc URL to your
Websphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
-Dserver.root="$WAS_HOME" $USER_INSTALL_PROP "$CLIENTSAS"
-classpath "$WAS_CLASSPATH";<list of your application jars and classes>
<fully qualified class name to run><your application parameters>
```

- Use the WebSphere Application Server launcher:

On Windows:

```
"%JAVA_HOME%\bin\java" "-Xbootclasspath/p:%WAS_BOOTCLASSPATH%"
-Djava.security.auth.login.config="%WAS_HOME%\properties\wsjaas_client.conf"
"-Dws.ext.dirs=<list of your application jars and classes;
%WAS_EXT_DIRS%;%WAS_USER_DIRS%">
-Djava.naming.provider.url=<an iiop URL or a corbaloc URL to your
Websphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
"-Dserver.root=%WAS_HOME%"
"%CLIENTSAS%" %USER_INSTALL_PROP% -classpath "%WAS_CLASSPATH%" com.ibm.ws.bootstrap
.WSLauncher
<fully qualified class name to run><your application parameters>
```

On Unix:

```
"$JAVA_HOME/bin/java" "-Xbootclasspath/p:$WAS_BOOTCLASSPATH"
-Djava.security.auth.login.config="$WAS_HOME/properties/wsjaas_client.conf"
"-Dws.ext.dirs=<list of your application jars and classes
>$WAS_EXT_DIRS;$WAS_USER_DIRS"
-Djava.naming.provider.url=<an iiop URL or a corbaloc URL to your
Websphere server machine name>
-Djava.naming.factory.initial=com.ibm.websphere.naming.WsnInitialContextFactory
"-Dserver.root=$WAS_HOME"
"$CLIENTSAS" $USER_INSTALL_PROP -classpath "$WAS_CLASSPATH" com.ibm.ws.bootstrap
.WSLauncher
<fully qualified class name to run><your application parameters>
```

For more information on iiop and corbaloc URLs, see “Developing applications that use JNDI” on page 1056.

### Usage scenario

View the Samples gallery for more information about application clients.

These samples do not include client applications that communicate with Enterprise server-side examples.

---

## Assembling Application Client Modules

### Before you begin

If you want to use existing J2EE 1.2 Web modules in your J2EE 1.3 application, migrate them to J2EE 1.3 first.

**Note:** This task only applies to J2EE application clients.

Assemble a client module to contain application client code. (Group enterprise beans, Web components, and resource adapter code in separate modules).

### Steps for this task

1. Start the Application Assembly Tool.
2. From the **New** tab, select **Application Client**. Click **OK**.  
The navigation tree now displays various sets of properties for configuring the new application client.
3. **(Optional)** Use the property dialog shown in the AAT workspace to change the default file name and location.
  - a. It is recommended that you change the display name so that it differs from the file name.
  - b. If you like, change the temporary location of the application client from the default location, `install_root/bin`.
4. Enter the main class filename and location.
  - a. Click **Browse** to locate the class file.
  - b. Select the archive containing the class files, and click **Select**.
  - c. Select the files you need from the archive, and click **OK**.
5. Define the assembly properties for the application client.
  - a. Right-click the assembly property in the navigation pane.
  - b. Select **New** in the right-click menu.
  - c. Enter values in the property dialog box.
  - d. Click **OK**.
6. Add files for the application client.
  - a. Right-click **Files** in the navigation pane.
  - b. Select **Add Files** from the right-click menu.
  - c. Locate the directory where the files are located and click **Select**.
  - d. Select the files to add and click **Add**.
  - e. In the **Selected files** window, click **OK**.
7. "Saving applications after assembly" on page 1161.

#### What to do next

Assemble other new modules of your choice, if needed:

- "Assembling EJB modules" on page 136
- "Assembling Resource Adapter modules" (not in this document).

You can also "Migrating application modules from J2EE 1.2 to J2EE 1.3" on page 1152.

Another option is to proceed directly to assembling the module. While assembling an application module, you can create any new modules that you need.

## **Application client modules**

An application client module is a JAR (Java ARchive) file containing a client for accessing a Java application.

The J2EE (Java™ 2 Platform for WebSphere Application Server Enterprise) Application Client and the Java Application Thin Client are supported by this product. Both clients provide first tier client support where Web applications are enabled with Java technologies.

The Java Application Thin Client provides users a lightweight, downloadable Java application run time that is capable of interacting with enterprise beans. This client

is designed to support those users who want a Java client application programming environment without the overhead of the J2EE platform on the client machine.

## Application client assembly settings

Use this page to specify assembly properties for J2EE application clients.

### File name (Required, String)

Specifies the file name of the application client module, relative to the top level of the Enterprise Archive (EAR) file.

If this is a stand-alone module, the file name is the full path name of the archive.

### Alternative DD

Specifies the file name for an alternative deployment descriptor file to use instead of the original deployment descriptor file in the module Java Archive (JAR) file.

This file is the postassembly version of the deployment descriptor file. (You can edit the original deployment descriptor file to resolve dependencies and security information. Directing the use of the alternative deployment descriptor allows you to keep the original deployment descriptor file intact). The value of the Alternative DD property must be the full path name of the deployment descriptor file relative to the module root directory. By convention, the file is in the ALT-INF directory. If this property is not specified, the deployment descriptor file is read directly from the module JAR file.

### Classpath

Specifies the full classpath containing the dependent code that is not contained in the application client JAR file.

Specify the values relative to the root of the EAR file and separate the values with spaces. Absolute values that reference files or directories on the hard drive are ignored. To specify classes that are not in JAR files but are in the root of the EAR file, use a period and forward slash (.). Consider the following example directory structure in which the file myapp.ear contains an application client JAR file named client.jar. Additional classes reside in class1.jar and class2.zip files. A class named xyz.class is not packaged in a JAR file but is in the root of the EAR file myapp.ear/client.jar myapp.ear/class1.jar myapp.ear/class2.zip myapp.ear/xyz.class. Specify class1.jar class2.zip ./ as the value of the Classpath property. (Name only the directory for .class files.)

### Display name (Required, String)

Specifies a short name that is intended for display by GUIs.

### Small icon

Specifies a JPEG or GIF file containing a small image (16x16 pixels).

The image is used as an icon to represent the application client in a GUI.

### Large icon

Specifies a JPEG or GIF file containing a large image (32x32 pixels).

The image is used as an icon to represent the application client in a GUI

### Description

Contains text describing the application client.

### **Main class (Required, String)**

Specifies the full path name of the main class for this application client.

### **WebSphere Application Server Enterprise Edition assembly settings for application clients**

Use this page to configure Enterprise Application Server functions for application clients.

**Own task:** A J2EE component's own task defines the task that can be used to identify a unit of work begun by this component to application profile-configured policies on downstream requests.

**Name:** The name of the task to be propagated on requests from this J2EE component.

Datatype String

**Description:** The description of the task.

An optional field provided as a convenience to the developer.

Datatype String

**Task references:** Task references enable developers to programmatically set the current task under which a component is executing.

**Name:** The name of the task reference. This name corresponds to the string that the developer uses to set the current task.

**Task:** The task that the run-time environment will associate with execution of the current component.

**Name:** The name of the task.

Programmatically setting the logical task name causes this task name to be associated with requests from this J2EE component.

**Description:** A description of the task.

An optional field provided as a convenience for the developer.

## **Environment entries assembly properties**

Use the assembly properties to change the settings for environment entries.

Environment entries define variables for use at run-time. The container makes the application environment entries available in a JNDI naming context (java:comp/env). An example environment entry is `maxExceptions`, describing the maximum number of tax exemptions that are allowed. The expected type of the environment entry is `java.lang.Integer` and its value is 15.

To view the settings for environment entries in the Application Assembly Tool, click on an application client (.jar file) in the navigation tree and then double-click on **Environment Entries** in the **Name** list.

**Name**

(Required) Specifies the name of the environment entry, relative to the java:comp/env context.

Data type String

**Value**

Specifies the value of the environment entry. The value must be a string that is valid for the constructor of the specified type that takes a single string parameter.

Data type String

**Type**

(Required) Specifies the Java type of the environment entry value that is expected by the module's code.

Valid values for this field include classes such as java.lang.Boolean, java.lang.String, java.lang.Integer, java.lang.Double, java.lang.Byte, java.lang.Short, java.lang.Long, and java.lang.Float.

|           |             |
|-----------|-------------|
| Data type | String      |
| Units     | Class names |
| Default   | String      |

**Description**

Describes the environment entries.

Data type String

## Deploying application clients on workstation platforms

### Before you begin

After developing an application client, deploy this application on client machines. *Deployment* consists of pulling together the various artifacts that the application client requires.

The *Application Client Resource Configuration Tool* (ACRCT) defines resources for the application client. These configurations are stored in the application client .ear file. The application client run time uses these configurations for resolving and creating an instance of the resources for the application client.

**Note:** This task only applies to J2EE application clients. Only perform this task if you configured your J2EE application client to use resource references.

### Steps for this task

1. Start the ACRCT and open an EAR file.
2. Configure new data source providers.
3. Configure mail providers and sessions.
4. Configure URL providers and sessions.
5. Configure Java messaging client resources.
6. Configure new environment entries.
7. (Optional) Remove application client resources.



8. "Saving applications after assembly" on page 1161

## Starting the Application Client Resource Configuration Tool and opening an EAR file

Before you begin

**Note:** This task only applies to J2EE application clients.

Steps for this task

1. Open a command prompt and change to the `install_root\bin` directory.
2. Run the `clientConfig.bat` file for a Windows system or the `clientConfig.sh` file for a UNIX system.
3. Open an EAR file within the Application Client Resource Configuration Tool (ACRCT):
  - Click **File > Open**.
  - Select the file and click **Open**.
4. Save your changes to the file and close the tool:
  - Click **File > Save**.
  - Click **File > Exit**.

## Data sources for application clients

The J2EE application client does not support looking up or directly accessing data source resources that have been configured on the WebSphere Application Server because the J2EE application client does not support Java 2 Connection Factories. To use a data source directly from the client application, you must configure your data source using the ACRCT. In addition, WebSphere Application Server and WebSphere Application Server clients do not provide client database drivers to be used directly from a J2EE application client. If your application client uses a database directly, you must provide the database drivers on the client machine. This can involve contacting your database vendor to acquire client database driver code and licenses. Instead of accessing the database directly, it is recommended that your client application use an enterprise bean. Accessing a database through an enterprise bean eliminates the need to have database drivers on the client machine, since the database access is handled by the enterprise bean running on the WebSphere Application Server. For a current list of providers that are supported on the WebSphere Application Server go the following site:

Supported Prerequisites page.

## Configuring new data source providers (JDBC providers) for application clients

Before you begin

During this task, you create new data source providers, also known as JDBC providers, for your application client. In a separate administrative task, install the Java code for the required data source provider on the client machine on which the application client resides.

Steps for this task

1. Start the tool and open the EAR file for which you want to configure the new data source provider. The EAR file contents display in a tree view.

2. Select the JAR file in which you want to configure the new data source provider from the tree.
3. Expand the JAR file to view its contents.
4. Click the **Data Source Providers** folder. Do one of the following:
  - Right-click the folder and click **New Provider**.
  - Click **Edit > New** on the menu bar.
5. Configure the data source provider properties in the resulting property dialog.
6. Click **OK** when you finish.
7. Click **File > Save** on the menu bar to save your changes.

### Configuring new data source providers

During this task, you will create new data source providers, also known as JDBC drivers, for your application client. In a separate administrative task, install the Java code for the required data source provider on the client machine where the application client resides.

#### Steps for this task

1. Start the ACRCT, click **File > Open**, and select the EAR file for which you want to configure the new data source provider. The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new data source provider from the tree.
3. Expand the JAR file to view its contents.
4. Right click the **Data Source Providers** folder and select **New Provider**.
5. Configure the data source provider properties in the resulting property dialog.
6. Click **OK**.
7. Click **File > Save** to save your changes.

### Example: Configuring data source provider and data source settings

The purpose of this article is to help you to configure data source provider and data source settings.

- Required fields:
  - Data Source Provider Properties page: name
  - Data Source Properties page: name, jndiName
- Special cases:
  - The user name and password fields have no equivalent xmi tags. You must specify these fields in the custom properties.
  - The password is encrypted when you use the Application Client Resource Configuration Tool (ACRCT). If you do not use the ACRCT, the field cannot be encrypted.
- Example (shown here on multiple lines for publication):

```
<resources.jdbc:JDBCProvider xmi:id="JDBCProvider_1" name="jdbcProvider:name"
  description="jdbcProvider:description" implementationClassName="jdbcProvider:ImplementationClass">
<classpath>jdbcProvider:classpath</classpath>
<factories xmi:type="resources.jdbc:WAS40DataSource" xmi:id="WAS40DataSource_1" name="jdbcFactory:name"
  jndiName="jdbcFactory:jndiName" description="jdbcFactory:description"
  databaseName="jdbcFactory:databaseName">
<propertySet xmi:id="J2EEResourcePropertySet_13">
<resourceProperties xmi:id="J2EEResourceProperty_13" name="jdbcFactory:customName"
  value="jdbcFactory:customValue"/>
<resourceProperties xmi:id="J2EEResourceProperty_14" name="user" value="jdbcFactory:user"/>
<resourceProperties xmi:id="J2EEResourceProperty_15" name="password"
```

```

    value="{xor}NTs9PBk+PCswLSZ1MT4y0g==" />
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_14">
<resourceProperties xmi:id="J2EEResourceProperty_16" name="jdbcProvider:customName"
    value="jdbcProvider:customeValue"/>
</propertySet>
</resources.jdbc:JDBCProvider>

```

## Data source provider settings for application clients

Use this page to create a data source under a JDBC provider which provides the specific JDBC driver implementation class.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Data Source Providers >** and click **New**. The following fields appear on the **General** tab:

**Name:** Specifies the display name for the data source.

For example you can set this field to *Test Data Source*.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** Specifies a text description for the resource.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Class Path:** A list of paths or jarfile names which together form the location for the resource provider classes.

**Implementation class:** Use this setting to perform database specific functions.

|           |   |
|-----------|---|
| Data type | String  |
| Default   | Dependent on JDBC driver implementation class |

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

## Data source properties for application clients

Use this page to create or modify the V5.0 data sources.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **Data Source Providers > Data source provider instance**. Right-click **Data Sources** and click **New**. The following fields appear on the **General** tab:

**Name:** Specifies the display name of this data source.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** Specifies a text description of the data source.

Data type String

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**Database Name:** The name of the database to which you want to connect.

**User:** Use the user ID with the Password property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the User ID property, you must also specify a value for the Password property. The connection factory User ID and Password properties are used if the calling application does not provide a userid and password explicitly.

**Password:** Use the password with the User ID property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the User ID property, you must also specify a value for the Password property.

**Re-Enter Password:** Confirms the password.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

## Configuring new data sources for application clients

### Before you begin

During this task, you create new data sources for your application client.

### Steps for this task

1. Click the data source provider for which you want to create a data source in the tree. Do one of the following:
  - Configure a new data source provider.
  - Click an existing data source provider.
2. Expand the data source provider to view its **Data Sources** folder.
3. Click the folder. Do one of the following:
  - Right-click the folder and click **New Factory**.
  - Click **Edit > New** on the menu bar.
4. Configure the data source properties in the resulting property dialog.
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

## Mail providers and mail sessions for the Application Client Assembly Tool

A JavaMail service provider is a driver that allows a JavaMail API to interact with mail servers running on a particular mail protocol. WebSphere Application Server includes service providers, also known as *protocol providers*, for mail protocols including Simple Mail Transfer Protocol (SMTP), Internet Message Access Protocol (IMAP), and Post Office Protocol 3 (POP3).

Mail provider encapsulates a collection of protocol providers. For example, WebSphere Application Server has a built-in mail provider that encompasses the three protocol providers: SMTP, IMAP and POP3. These protocol providers are installed as the default and are sufficient for most applications.

If you have a particular application that requires custom protocol providers, you must first follow the steps outlined in "JavaMail API Design Specification, V1.2, Chapter 5 - The Mail Session" to install your own protocol providers. See JavaMail: Resources for learning, for a link to this documentation.

Ensure every mail session is defined under a parent mail provider. Select a mail provider first and then create your new mail session.

## Configuring mail providers and sessions for application clients

### Before you begin

Use the Application Client Resource Configuration Tool (ACRCT) to edit the configurations of JavaMail sessions and providers for your application clients to use.

### Steps for this task

1. Open the ACRCT.
2. Open an EAR file.
3. Locate the JavaMail objects in the tree that displays.

For example, if your file contains JavaMail sessions, expand **Resources** > *application.jar* > **JavaMail Providers** > *java\_mail\_provider\_instance* > **JavaMail Sessions**.

In this example, *java\_mail\_provider\_instance* is a particular JavaMail provider.

### Results

The JavaMail session instances are located in the **JavaMail Sessions** folder.

### Mail provider settings for application clients

Use this page to implement the JavaMail API and create mail sessions.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File** > **Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Mail Providers** > and click **New**. The following fields appear on the **General** tab:

**Name:** The name of the JavaMail resource provider.

**Description:** An optional description for the resource provider.

**Class Path:** Specifies a list of paths or JAR file names which together form the location for the resource provider classes.

**Protocol:** Specifies the name of the protocol.

**Classname:** Specifies the name of the class implementing the protocol. Leave this field blank if you want to use the default implementation.

**Type:** This menu contains the following two values: TRANSPORT or STORE.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### **Mail session settings for application clients**

Use this page to configure mail session properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **Mail Providers > mail provider instance**. Right-click **Mail Sessions** and click **New**. The following fields appear on the **General** tab:

**Name:** Represents the administrative name of the JavaMail session object.

**Description:** Provides an optional description for your administrative records.

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**Mail Transport Host:** Specifies the server to connect to when sending mail.

**Mail Transport Protocol:** Specifies the transport protocol to use when sending mail.

**Mail Transport User:** Specifies the user ID to use when the mail transport host requires authentication.

**Mail Transport Password:** Specifies the password to use when the mail transport host requires authentication.

**Re-Enter Password:** Confirms the password.

**Mail From:** Specifies the mail originator.

**Mail Store Host:** Mail account host (or "domain") name.

**Mail Store User:** The user ID of the mail account.

**Mail Store Password:** The password of the mail account.

**Re-Enter Password:** Confirms the password.

**Mail Store Protocol:** Specifies the protocol to be used when receiving mail.

**Mail Debug:** When true, JavaMail interaction with mail servers, along with these mail session properties will be printed to stdout.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### **Example: Configuring JavaMail provider and JavaMail session settings for application clients**

The purpose of this article is to help you configure JavaMail provider and JavaMail session settings.

- Required fields:
  - JavaMail Provider Properties page: name, and at least one protocol provider
  - JavaMail Session Properties page: name, jndiName, mail transport protocol, mail store protocol
- Special cases:
  - The password is encrypted when using the ACRCT tool. Without the tool, you cannot encrypt this field.
- Example:

```
<resources.mail:MailProvider xmi:id="MailProvider_1" name="Default Mail Provider"
description="IBM JavaMail Implementation">
<classpath>mailProvider:classpath</classpath>
<factories xmi:type="resources.mail:MailSession" xmi:id="MailSession_1" name=
"mailSession:name" jndiName="mailSession:jndiName" description="mailSession:description"
mailTransportHost="mailSession:mailTransportHost" mailTransportUser=
"mailSession:mailTransportUser" mailTransportPassword="{xor}Mj42Mww6LCw2MDF1MT4y0g=="
mailFrom="mailSession:mailFrom" mailStoreHost="mailSession:mailStoreHost" mailStoreUser=
"mailSession:mailStoreUser" mailStorePassword="{xor}Mj42Mww6LCw2MDF1MT4y0g==" debug="true"
mailTransportProtocol="ProtocolProvider_1" mailStoreProvider="ProtocolProvider_1">
<propertySet xmi:id="J2EEResourcePropertySet_1">
<resourceProperties xmi:id="J2EEResourceProperty_1" name="mailSession:customName" value=
"mailSession:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_2">
<resourceProperties xmi:id="J2EEResourceProperty_2" name="mailProvider:customName" value=
"mailProvider:customValue"/>
</propertySet>
<protocolProviders xmi:id="ProtocolProvider_1" protocol="smtp" classname="smtp:className"
/>
<protocolProviders xmi:id="ProtocolProvider_2" protocol="pop3" classname="pop3:className"
/>
<protocolProviders xmi:id="ProtocolProvider_3" protocol="imap" classname="imap:className"
/>
</resources.mail:MailProvider>
```

## **Configuring new mail sessions for application clients**

### Before you begin

During this task, you configure new mail sessions for your application client. The mail sessions are associated with the preconfigured default mail provider supplied by the product.

### Steps for this task

1. Start the tool and open the EAR file for the new JavaMail session.  
The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new JavaMail session.
3. Expand the JAR file to view its contents.
4. Click **JavaMail Providers > MailProvider > JavaMail Sessions**. Do one of the following:
  - Right-click the **JavaMail Sessions** folder and select **New Factory**.
  - Click **Edit > New** on the menu bar.
5. Configure the JavaMail session properties in the resulting property dialog.
6. Click **OK** when you finish.
7. Click **File > Save** on the menu bar to save your changes.

## URLs for application clients

A *Uniform Resource Locator* (URL) is an identifier that points to an electronically accessible resource, such as a directory file on a machine in a network, or a document stored in a database.

URLs appear in the format *scheme:scheme\_information*.

You can represent a *scheme* as *http*, *ftp*, *file*, or another term that identifies the type of resource and the mechanism by which you can access the resource.

In a World Wide Web browser location or address box, a URL for a file available using HyperText Transfer Protocol (HTTP) starts with *http:*. An example is *http://www.ibm.com*. Files available using File Transfer Protocol (FTP) start with *ftp:*. Files available locally start with *file:*.

The *scheme\_information* commonly identifies the Internet machine making a resource available, the path to that resource, and the resource name. The *scheme\_information* for HTTP, FTP and File generally starts with two slashes (*//*), then provides the Internet address separated from the resource path name with one slash (*/*). For example,

*http://www-4.ibm.com/software/webservers/appserv/library.html*.

For HTTP and FTP, the path name ends in a slash when the URL points to a directory. In such cases, the server generally returns the default index for the directory.

## URL providers for the Application Client Resource Configuration Tool

A URL provider implements the functionality for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP). This provider, comprised of a pair of classes, extends the *java.net.URLStreamHandler* and *java.net.URLConnection* classes.

## Configuring new URL providers for application clients

Before you begin



During this task, you create URL providers and URLs for your client application. In a separate administrative task, you must install the Java code for the required URL provider on the client machine on which the client application resides.

#### Steps for this task

1. Start the tool and open the EAR file for which you want to configure the new URL provider. The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new URL provider from the tree.
3. Expand the JAR file to view its contents.
4. Click the folder called **URL Providers**. Do one of the following:
  - Right-click the folder and click **New Provider**.
  - Click **Edit -> New** on the menu bar.
5. Configure the URL provider properties in the resulting property dialog.
6. Click **OK** when you finish.
7. Click **File -> Save** on the menu bar to save your changes.

### **Configuring URL providers and sessions using the Application Client Resource Configuration Tool**

#### Before you begin

Use the Application Client Resource Configuration Tool (ACRCT) to edit the configurations of URL providers and URLs to be used by your application clients.

#### Steps for this task

1. Open the ACRCT.
2. Open an EAR file.
3. Locate the URL objects in the tree that displays.

For example, if your file contains URL providers and URLs, expand **Resources -> application.jar -> URL Providers -> url\_provider\_instance** where *url\_provider\_instance* is a particular URL provider.
4. If you expand the tree further, you will also see the **URLs** folders containing the URL instances for each URL provider instance.

### **URL settings for application clients**

Use this page to implement the functionality for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP).

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **URL Providers** > *URL provider instance*. Right-click **URLs** and click **New**. The following fields appear on the **General** tab.

This provider, comprised of classes, extends the `java.net.URLStreamHandler` and `java.net.URLConnection` classes.

**Name:** Administrative name for the URL

**Description:** Optional description of the URL, for your administrative records

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**URL:** A Uniform Resource Locator (URL) name that points to an internet or intranet resource. For example: `http://www.ibm.com`

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### **URL provider settings for application clients**

Use this page create new URLs..

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **URL Providers >** and click **New**. The following fields appear on the **General** tab.

A URL provider implements the functionality for a particular URL protocol, such as Hyper Text Transfer Protocol (HTTP). This provider, comprised of classes, extends the `java.net.URLStreamHandler` and `java.net.URLConnection` classes.

**Name:** Administrative name for the URL

**Description:** Optional description of the URL, for your administrative records

**Class Path:** A list of paths or jarfile names which together form the location for the resource provider classes.

**Protocol:** Protocol supported by this stream handler. For example, "nntp", "smtp", "ftp", etc.

To use the default protocol, leave this field blank.

**Stream handler class:** Fully qualified name of a User-defined Java class that extends `java.net.URLStreamHandler` for a particular URL protocol, such as ftp.

To use the default stream handler, leave this field blank.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### **Example: Configuring URL and URL provider settings for application clients**

The purpose of this article is to help you to configure URL and URL provider settings.

- Required fields:
  - URL Properties page: name, jndiName, url
  - URL Provider Properties page: name
- Example:

```
<resources.url:URLProvider xmi:id="URLProvider_1" name="urlProvider:name" description=
"urlProvider:description" streamHandlerClassName="urlProvider:streamHandlerClass" protocol=
"urlProvider:protocol">
<classpath>urlProvider:classpath</classpath>
<factories xmi:type="resources.url:URL" xmi:id="URL_1" name="urlFactory:name" jndiName=
"urlFactory:jndiName" description="urlFactory:description" spec="urlFactory:url">
<propertySet xmi:id="J2EEResourcePropertySet_18">
<resourceProperties xmi:id="J2EEResourceProperty_20" name="urlFactory:customName" value=
"urlFactory:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_19">
<resourceProperties xmi:id="J2EEResourceProperty_21" name="urlProvider:customName" value=
"urlProvider:customValue"/>
</propertySet>
</resources.url:URLProvider>
```

## Configuring new URLs with the Application Client Resource Configuration Tool

### Before you begin

During this task, you create URLs for your client application.

### Steps for this task

1. Click the URL provider for which you want to create a URL in the tree. Do one of the following:
  - Configure a new URL provider.
  - Click an existing URL provider.
2. Expand the URL provider to view the **URLs** folder.
3. Click the folder. Do one of the following:
  - Right-click the folder and click **New Factory**.
  - Click **Edit -> New** on the menu bar.
4. Configure the URL properties in the resulting property dialog.
5. Click **OK** when you finish.
6. Click **File -> Save** in the menu bar to save your changes.

## WebSphere asynchronous messaging using the Java Message Service API for the Application Client Resource Configuration Tool

WebSphere Application Server supports asynchronous messaging as a method of communication based on the Java Message Service (JMS) programming interface. The JMS interface provides a common way for Java programs (clients and J2EE applications) to create, send, receive, and read asynchronous requests, as JMS messages.

This topic provides an overview of asynchronous messaging using JMS support provided by the WebSphere Application Server.

The base support for asynchronous messaging using the JMS API provides the common set of JMS interfaces and associated semantics that define how a JMS client can access the facilities of a JMS provider. This support enables WebSphere J2EE applications, as JMS clients, to exchange messages asynchronously with other JMS clients, by using JMS destinations (queues or topics). An J2EE application can use JMS queue destinations for point-to-point messaging and JMS topic destinations for Pub and Sub messaging. A J2EE application can explicitly poll for messages on a destination then retrieve messages for processing by business logic beans (enterprise beans).

With the base JMS/XA support, the J2EE application uses standard JMS calls to process messages, including any responses or outbound messaging. An enterprise bean can handle responses acting as a sender bean, or within the enterprise bean that receives the incoming messages. Optionally, this process can use two-phase commit within the scope of a transaction. This level of functionality for asynchronous messaging is called *bean-managed messaging*, and gives an enterprise bean complete control over the messaging infrastructure; for example, connection and session pool management. The common container has no role in *bean-managed messaging*.

WebSphere Application Server also supports automatic asynchronous messaging using message-driven beans (a type of enterprise bean defined in the EJB 2.0 specification) and JMS listeners (part of the JMS application server facilities). Messages are automatically retrieved from JMS destinations, optionally within a transaction, then sent to the message-driven bean in a J2EE application, without the application having to explicitly poll JMS destinations.

## Configuring Java messaging client resources

### Before you begin

In a separate administrative task, install the Java Message Service (JMS) client on the client machine where the application client resides. The messaging product vendor must provide an implementation of the JMS client. For more information, see your messaging product documentation.

During this task, you create new JMS provider configurations for your application client. The application client can use a messaging service through the Java Message Service APIs. A JMS provider provides two kinds of J2EE factories. One is a *JMS connection factory*, and the other is a *JMS destination factory*.

### Steps for this task

1. Start the ACRCT, click **File > Open**, and select the EAR file for which you want to configure the new JMS provider. The EAR file contents display in a tree view.
2. Select the JAR file in which you want to configure the new JMS provider from the tree.
3. Expand the JAR file to view its contents.
4. Click the **JMS Providers** folder and click **New Provider**.
5. Configure the JMS provider properties in the resulting property dialog.
6. Click **OK**.
7. Click **File > Save**.

## Configuring new JMS providers with the Application Client Resource Configuration Tool

### Before you begin

During this task, you will create new JMS provider configurations for your application client. The application client can make use of a messaging service through the Java Message Service APIs. A JMS provider provides two kinds of J2EE factories. One is a JMS Connection factory, and the other is a JMS destination factory.

In a separate administrative task, you must install the JMS client on the client machine where the application client resides. The messaging product vendor must provide an implementation of the JMS client. For more information, see your messaging product documentation.

### Steps for this task

1. Start the tool and open the EAR file for which you want to configure the new JMS provider. The EAR file contents will be displayed in a tree view.
2. From the tree, select the JAR file in which you want to configure the new JMS provider.
3. Expand the JAR file to view its contents.
4. Click the folder called **JMS Providers**. Do one of the following:
  - Right-click the folder and select **New Provider**.
  - On the menu bar, click **Edit -> New**.
5. In the resulting property dialog, configure the JMS provider properties.
6. When finished, click **OK**.
7. On the menu bar, click **File -> Save** to save your changes.

### **JMS provider settings for application clients**

Use this page to configure properties of the JMS provider, if you want to use a JMS provider other than the internal WebSphere JMS provider or the MQSeries JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **JMS Providers > click New**. The following fields appear on the **General** tab.

**Name:** The name by which the JMS provider is known for administrative purposes.

Data type String

**Description:** A description of the JMS provider, for administrative purposes

Data type String

**Class Path:** A list of paths or jarfile names which together form the location for the resource provider classes.

**Context factory class:** The Java classname of the initial context factory for the JMS provider.

For example, for an LDAP service provider the value has the form:  
com.sun.jndi.ldap.LdapCtxFactory.

Data type String

**Provider URL:** The JMS provider URL for external JNDI lookups.

For example, an LDAP URL for a JMS provider has the form:  
ldap://hostname.company.com/contextName.

Data type String

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### WebSphere queue connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected queue connection factory for use with the internal WebSphere JMS provider that is installed with WebSphere Application Server. These configuration properties control how connections are created to the associated JMS queue destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **WAS Queue Connection Factories** and click **New**. The following fields appear on the **General** tab.

A queue connection factory is used to create JMS connections to queue destinations. The queue connection factory is created by the internal WebSphere JMS provider. A queue connection factory for the internal WebSphere JMS provider has the following properties:

**Name:** The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

Data type String

**Description:** A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String  
Default Null

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**User:** The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the userid and password to the JMS server.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Password:** The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Re-Enter Password:** Confirms the password.

**Node:** The WebSphere node name of the administrative node where the JMS server runs for this connection factory. Connections created by this factory connect to that JMS server.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Application Server:** Enter the name of the application server. This name is not the host name of the machine, but the name of the configured application server.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### **WebSphere topic connection factory settings for application clients**

Use this panel to view or change the configuration properties of the selected topic connection factory for use with the internal WebSphere JMS provider. These configuration properties control how connections are created to the associated JMS topic destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected

JAR file > **JMS Providers** > *JMS provider instance*. Right-click **WAS Topic Connection Factories** and click **New**. The following fields appear on the **General** tab.

A topic connection factory is used to create JMS connections to topic destinations. The topic connection factory is created by the associated JMS provider. A topic connection factory for the internal WebSphere JMS provider has the following properties.

**Name:** The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** A description of this topic connection factory for administrative purposes within IBM WebSphere Application Server.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**User:** The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method createTopicConnection(). The JMS client flows the userid and password to the JMS server.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Password:** The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Re-Enter Password:** Confirms the password.

**Node:** The WebSphere node name of the administrative node where the JMS server runs for this connection factory. Connections created by this factory connect to that JMS server.



|           |   |
|-----------|---|
| Data type | Enum  |
| Default   | Null  |
| Range     | Pull-down list of nodes in the WebSphere administrative domain. |

**Application Server:** Enter the name of the application server. This name is not the host name of the machine, but the name of the configured application server.

**Port:** Which of the two ports that connections use to connect to the JMS Server. The QUEUED port is for full-function JMS publish/subscribe support, the DIRECT port is for non-persistent, non-transactional, non-durable subscriptions only.

**Note:** Message-driven beans cannot use the direct listener port for publish/subscribe support. Therefore, any topic connection factory configured with **Port** set to Direct cannot be used with message-driven beans.

|           |  |
|-----------|--|
| Data type | Enum   |
| Units     | Not applicable   |
| Default   | QUEUED   |
| Range     | <p><b>QUEUED</b><br/>The listener port used for full-function JMS-compliant, publish/subscribe support.</p> <p><b>DIRECT</b><br/>The listener port used for direct TCP/IP connection (non-transactional, non-persistent, and non-durable subscriptions only) for publish/subscribe support.</p> <p>The TCP/IP port numbers for these ports are defined on the WebSphere Internal JMS Server.</p> |

**Client Id:** The JMS client identifier used for connections to the MQSeries queue manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### WebSphere queue destination settings for application clients

Use this panel to view or change the configuration properties of the selected queue destination for use with the WebSphere JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **WAS Queue Destinations** and click **New**. The following fields appear on the **General** tab.

A queue destination is used to configure the properties of a JMS queue. Connections to the queue are created by the associated queue connection factory

for the internal WebSphere JMS provider. A queue for use with the internal WebSphere JMS provider has the following properties.

**Name:** The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** A description of the queue, for administrative purposes

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**Persistence:** Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

|           |                     |
|-----------|---------------------|
| Data type | Enum                |
| Units     | Not applicable      |
| Default   | APPLICATION_DEFINED |
| Range     |                     |

**Application defined**

Messages on the destination have their persistence defined by the application that put them onto the queue.

**Queue defined**

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

**Persistent**

Messages on the destination are persistent.

**Non persistent**

Messages on the destination are not persistent.

**Priority:** Whether the message priority for this destination is defined by the application or the **Specified priority** property

|           |                     |
|-----------|---------------------|
| Data type | Enum                |
| Units     | Not applicable      |
| Default   | APPLICATION_DEFINED |

Range

**Application defined**

The priority of messages on this destination is defined by the application that put them onto the destination.

**Queue defined**

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

**Specified**

The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

**Specified Priority:** If the **Priority** property is set to **Specified**, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to **Specified**, messages sent to this queue have the priority value specified by this property.

|           |  |
|-----------|--|
| Data type | Integer  |
| Units     | Message priority level                           |
| Default   | Null   |
| Range     | 0 (lowest priority) through 9 (highest priority) |

**Expiry:** Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

|           |                     |
|-----------|---------------------|
| Data type | Enum                |
| Units     | Not applicable      |
| Default   | APPLICATION_DEFINED |
| Range     |                     |

**Application defined**

The expiry timeout for messages on this queue is defined by the application that put them onto the queue.

**Specified**

The expiry timeout for messages on this queue is defined by the **Specified expiry** property. *If you select this option, you must define a timeout on the **Specified expiry** property.*

**Unlimited**

Messages on this queue have no expiry timeout, so those messages never expire.

**Specified Expiry:** If the **Expiry timeout** property is set to **Specified**, type here the number of milliseconds (greater than 0) after which messages on this queue expire

|           |              |
|-----------|--------------|
| Data type | Integer      |
| Units     | Milliseconds |
| Default   | Null         |

- Range      Greater than or equal to 0
- 0 indicates that messages never timeout
  - Other values are an integer number of milliseconds

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### WebSphere topic destination settings for application clients

Use this panel to view or change the configuration properties of the selected topic destination for use with the internal WebSphere JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers** > *JMS provider instance*. Right-click **WAS Topic Destinations** and click **New**. The following fields appear on the **General** tab.

A topic destination is used to configure the properties of a JMS topic for the associated JMS provider. Connections to the topic are created by the associated topic connection factory. A topic for use with the internal WebSphere JMS provider has the following properties.

**Name:** The name by which the topic is known for administrative purposes.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** A description of the topic, for administrative purposes within IBM WebSphere Application Server.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**Topic Name:** The name of the topic as defined to the JMS provider.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Persistence:** Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

|           |                     |
|-----------|---------------------|
| Data type | Enum                |
| Units     | Not applicable      |
| Default   | APPLICATION_DEFINED |

Range

**Application defined**

Messages on the destination have their persistence defined by the application that put them onto the queue.

**Queue defined**

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

**Persistent**

Messages on the destination are persistent.

**Non persistent**

Messages on the destination are not persistent.

**Priority:** Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type

Enum

Units

Not applicable

Default

APPLICATION\_DEFINED

Range

**Application defined**

The priority of messages on this destination is defined by the application that put them onto the destination.

**Queue defined**

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

**Specified**

The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

**Specified Priority:** If the **Priority** property is set to **Specified**, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to **Specified**, messages sent to this queue have the priority value specified by this property.

Data type

Integer

Units

Message priority level

Default

Null

Range

0 (lowest priority) through 9 (highest priority)

**Expiry:** Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

Data type

Enum

Units

Not applicable

Default

APPLICATION\_DEFINED

Range

#### Application defined

The expiry timeout for messages on this queue is defined by the application that put them onto the queue.

#### Specified

The expiry timeout for messages on this queue is defined by the **Specified expiry** property. *If you select this option, you must define a timeout on the **Specified expiry** property.*

#### Unlimited

Messages on this queue have no expiry timeout, so those messages never expire.

**Specified Expiry:** If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | Milliseconds  |
| Default   | Null  |
| Range     | Greater than or equal to 0 <ul style="list-style-type: none"><li>• 0 indicates that messages never timeout</li><li>• Other values are an integer number of milliseconds</li></ul> |

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### MQSeries queue connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected queue connection factory for use with the MQSeries JMS provider. These configuration properties control how connections are created to the associated JMS queue destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **MQ Queue Connection Factories** and click **New**. The following fields appear on the **General** tab.

A queue connection factory is used to create JMS connections to queue destinations. The queue connection factory is created by the MQSeries JMS provider. A queue connection factory for the MQSeries JMS provider has the following properties.

#### Note:

- The property values that you specify must match the values that you specified when configuring MQSeries for JMS resources. For more information about configuring MQSeries JMS resources, see the MQSeries *Using Java* book.

- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

**Name:** The name by which this queue connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS connection factories across the WebSphere administrative domain.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**User:** The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the userid and password to the JMS server.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Password:** The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Re-Enter Password:** Confirms the password.

**Queue Manager:** The name of the MQSeries queue manager for this connection factory.

Connections created by this factory connect to that queue manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Host:** The name of the host on which the WebSphere MQ queue manager runs, for client connection only.

|           |                         |
|-----------|-------------------------|
| Data type | String                  |
| Default   | Null                    |
| Range     | A valid TCP/IP hostname |

**Port:** The TCP/IP port number used for connection to the WebSphere MQ queue manager, for client connection only.

This port must be configured on the WebSphere MQ queue manager.

|           |   |
|-----------|---|
| Data type | Integer   |
| Default   | Null  |
| Range     | A valid TCP/IP port number, configured on the WebSphere MQ queue manager. |

**Channel:** The name of the channel used for connection to the WebSphere MQ queue manager, for client connection only.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Default   | Null                          |
| Range     | 1 through 20 ASCII characters |

**Transport type:** Specifies whether the WebSphere MQ client connection or JNI bindings are used for connection to the WebSphere MQ queue manager. The external JMS provider controls the communication protocols between JMS clients and JMS servers. Tune the transport type when you are using non-ASF nonpersistent, nondurable, nontransactional messaging or when you want to satisfy security issues and the client is local to the queue manager node.

**Data type**  
**Units**  
**Default**  
**Range**

**Enum**  
Not applicable  
BINDINGS

**BINDINGS**

JNI bindings are used to connect to the queue manager. BINDINGS is a shared memory protocol and can only be used when the queue manager is on the same node as the JMS client and comes at some security risks that should be addressed through the use of EJB roles.

**CLIENT**

WebSphere MQ client connection is used to connect to the queue manager. CLIENT is a typical TCP-based protocol.

**DIRECT**

For WebSphere MQ Event Broker using DIRECT mode. DIRECT is a lightweight sockets protocol used in nontransactional, nondurable and nonpersistent Publish/Subscribe messaging. DIRECT is only works for clients and message-driven beans using the non-ASF protocol.

**QUEUED**

QUEUED is a standard TCP protocol.



Data type  
Recommended

Enum

**Queue connection factory transport type**

BINDINGS is faster by 30% or more, but it lacks security. When you have security concerns, BINDINGS is more desirable than CLIENT.

**Topic connection factory transport type**

DIRECT is the fastest and should be used where possible. Use BINDINGS when you want to satisfy additional security tasks and the queue manager is local to the JMS client. QUEUED is fallback for all other cases. Note, WebSphere MQ 5.3 before CSD2 with the DIRECT setting can lose messages when used with message-driven beans and under load. This also happens with client-side based applications unless the broker's maxClientQueueSize is set to 0. You can set this to 0 with the command `#wempschangeproperties WAS_nodeName_server1 -e default -o DynamicSubscriptionEngine -n maxClientQueueSize -v 0 -x executionGroupUUID`, where executionGroupUUID can be found by starting the broker and looking in the Event Log/Applications for event 2201. This value is usually ffffffff-0000-0000-000000000000.

**Client ID:** The JMS client identifier used for connections to the MQSeries queue manager.

Data type String

**CCSID:** The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

Data type String

For more information about supported CCSIDs, and about converting between message data from one coded character set to another, see the *WebSphere MQ System Administration* and the *WebSphere MQ Application Programming Reference* books. These are available from the WebSphere MQ messaging multiplatform and platform-specific books Web pages, the IBM Publications Center (<http://www.ibm.com/publications>), or from the WebSphere MQ collection kit, SK2T-0730.

**Message Retention:** Select this tick box to specify that unwanted messages are to be left on the queue. Otherwise, unwanted messages are dealt with according to their disposition options.

|           |  |
|-----------|--|
| Data type | Enum   |
| Units     | Not applicable   |
| Default   | Cleared  |
| Range     | <b>Selected</b><br>Unwanted messages are left on the queue.                                |
|           | <b>Cleared</b><br>Unwanted messages are dealt with according to their disposition options. |

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### MQSeries topic connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected topic connection factory for use with the MQSeries JMS provider. These configuration properties control how connections are created to the associated JMS topic destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **MQ Topic Connection Factories** and click **New**. The following fields appear on the **General** tab.

A topic connection factory is used to create JMS connections to topic destinations. The topic connection factory is created by the MQSeries JMS provider. A topic connection factory for the MQSeries JMS provider has the following properties.

**Note:**

- The property values that you specify must match the values that you specified when configuring MQSeries JMS resources. For more information about configuring MQSeries JMS resources, see the MQSeries *Using Java* book.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

**Name:** The name by which this topic connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the JMS provider.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** A description of this topic connection factory for administrative purposes within IBM WebSphere Application Server.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**JNDI Name:** The JNDI name that is used to bind the topic connection factory into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Default   | Null                          |
| Range     | 1 through 45 ASCII characters |

**User:** The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User** property, you must also specify a value for the **Password** property.

The connection factory **User** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createTopicConnection()`. The JMS client flows the userid and password to the JMS server.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Password:** The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**Re-Enter Password:** Confirms the password.

**Queue Manager:** The name of the MQSeries queue manager for this connection factory. Connections created by this factory connect to that queue manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Host:** The name of the host on which the WebSphere MQ queue manager runs, for client connection only.

|           |                         |
|-----------|-------------------------|
| Data type | String                  |
| Default   | Null                    |
| Range     | A valid TCP/IP hostname |

**Port:** The TCP/IP port number used for connection to the WebSphere MQ queue manager, for client connection only.

This port must be configured on the WebSphere MQ queue manager.

|           |   |
|-----------|---|
| Data type | Integer   |
| Default   | Null  |
| Range     | A valid TCP/IP port number, configured on the WebSphere MQ queue manager. |

**Channel:** The name of the channel used for connection to the WebSphere MQ queue manager, for client connection only.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Default   | Null                          |
| Range     | 1 through 20 ASCII characters |

**Transport Type:** Whether MQSeries client connection or JNDI bindings is used for connection to the MQSeries queue manager.

|           |  |
|-----------|--|
| Data type | Enum   |
| Units     | Not applicable   |
| Default   | BINDINGS   |
| Range     | <b>CLIENT</b><br>MQSeries client connection is used to connect to the MQSeries queue manager.<br><b>BINDINGS</b><br>JNDI bindings are used to connect to the MQSeries queue manager. |

**Client Id:** The JMS client identifier used for connections to the MQSeries queue manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**CCSID:** The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Broker Control Queue:** The name of the broker control queue, to which all command messages (except publications and requests to delete publications) are sent

The name of the broker control queue. Publisher and subscriber applications, and other brokers, send all command messages (except publications and requests to delete publications) to this queue.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Default   | Null                          |
| Range     | 1 through 48 ASCII characters |

**Broker Queue Manager:** The name of the MQSeries queue manager that provides the Pub/Sub message broker.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Default   | Null                          |
| Range     | 1 through 48 ASCII characters |

**Broker Pub Queue:** The name of the broker's input queue that receives all publication messages for the default stream

The name of the broker's input queue (stream queue) that receives all publication messages for the default stream. Applications can also send requests to delete publications on the default stream to this queue.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Default   | Null                          |
| Range     | 1 through 48 ASCII characters |

**Broker Sub Queue:** The name of the broker queue from which non-durable subscription messages are retrieved

The name of the broker's queue from which non-durable subscription messages are retrieved. The subscriber specifies the name of the queue when it registers a subscription.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Default   | Null                          |
| Range     | 1 through 48 ASCII characters |

**Broker CCSubQ:** The name of the broker's queue from which non-durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

The name of the broker queue from which non-durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Default   | Null                          |
| Range     | 1 through 48 ASCII characters |

**Broker Version:** Whether the message broker is provided by the MQSeries MA0C Supportpac or newer versions of WebSphere message broker products

|           |                |
|-----------|----------------|
| Data type | Enum           |
| Units     | Not applicable |
| Default   | Advanced       |

Range

**Advanced**

The message broker is provided by newer versions of WebSphere message broker products (MQ Integrator and MQ Publish and Subscribe)

**Basic**

The message broker is provided by the MQSeries MA0C SupportPac (MQSeries - Publish/Subscribe)

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### MQSeries queue destination settings for application clients

Use this panel to view or change the configuration properties of the selected queue destination for use with the MQSeries JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > JMS provider instance**. Right-click **MQ Queue Destinations** and click **New**. The following fields appear on the **General** tab.

A queue destination is used to configure the properties of a JMS queue. Connections to the queue are created by the associated queue connection factory for the MQSeries JMS provider. A queue for use with the MQSeries JMS provider has the following properties.

**Note:**

- The property values that you specify must match the values that you specified when configuring MQSeries JMS resources. For more information about configuring MQSeries JMS resources, see the MQSeries *Using Java* book.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

**Name:** The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Description:** A description of the queue, for administrative purposes

|           |        |
|-----------|--------|
| Data type | String |
| Default   | Null   |

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**Persistence:** Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

|           |  |
|-----------|--|
| Data type | Enum   |
| Units     | Not applicable   |
| Default   | APPLICATION_DEFINED  |
| Range     | <b>Application defined</b><br>Messages on the destination have their persistence defined by the application that put them onto the queue.<br><b>Queue defined</b><br>[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.<br><b>Persistent</b><br>Messages on the destination are persistent.<br><b>Non persistent</b><br>Messages on the destination are not persistent. |

**Priority:** Whether the message priority for this destination is defined by the application or the **Specified priority** property

|           |   |
|-----------|---|
| Data type | Enum  |
| Units     | Not applicable  |
| Default   | APPLICATION_DEFINED   |
| Range     | <b>Application defined</b><br>The priority of messages on this destination is defined by the application that put them onto the destination.<br><b>Queue defined</b><br>[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.<br><b>Specified</b><br>The priority of messages on this destination is defined by the <b>Specified priority</b> property. <i>If you select this option, you must define a priority on the <b>Specified priority</b> property.</i> |

**Specified Priority:** If the **Priority** property is set to Specified, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to Specified, messages sent to this queue have the priority value specified by this property.

|           |  |
|-----------|--|
| Data type | Integer  |
| Units     | Message priority level                           |
| Default   | Null   |
| Range     | 0 (lowest priority) through 9 (highest priority) |

**Expiry:** Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

|           |   |
|-----------|---|
| Data type | Enum  |
| Units     | Not applicable  |
| Default   | APPLICATION_DEFINED   |
| Range     | <p><b>Application defined</b><br/>The expiry timeout for messages on this queue is defined by the application that put them onto the queue.</p> <p><b>Specified</b><br/>The expiry timeout for messages on this queue is defined by the <b>Specified expiry</b> property. <i>If you select this option, you must define a timeout on the <b>Specified expiry</b> property.</i></p> <p><b>Unlimited</b><br/>Messages on this queue have no expiry timeout, so those messages never expire.</p> |

**Specified Expiry:** If the **Expiry timeout** property is set to Specified, type here the number of milliseconds (greater than 0) after which messages on this queue expire

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | Milliseconds  |
| Default   | Null  |
| Range     | <p>Greater than or equal to 0</p> <ul style="list-style-type: none"> <li>• 0 indicates that messages never timeout</li> <li>• Other values are an integer number of milliseconds</li> </ul> |

**Base Queue Name:** The name of the queue to which messages are sent, on the queue manager specified by the **Base queue manager name** property

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Base Queue Manager Name:** The name of the MQSeries queue manager to which messages are sent

This queue manager provides the queue specified by the **Base queue name** property.

|           |   |
|-----------|---|
| Data type | String  |
| Units     | En_US ASCII characters  |
| Default   | Null  |
| Range     | A valid MQSeries Queue Manager name, as 1 through 48 ASCII characters |

**CCSID:** The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Integer encoding:** If native encoding is not enabled, select whether integer encoding is normal or reversed.



|           |  |
|-----------|--|
| Data type | Enum   |
| Units     | Not applicable   |
| Default   | NORMAL   |
| Range     | <b>NORMAL</b><br>Normal integer encoding is used.  |
|           | <b>REVERSED</b><br>Reversed integer encoding is used.  |
|           | For more information about encoding properties, see the WebSphere MQ <i>Using Java</i> document. |

**Decimal encoding:** If native encoding is not enabled, select whether decimal encoding is normal or reversed.

|           |  |
|-----------|--|
| Data type | Enum   |
| Units     | Not applicable   |
| Default   | NORMAL   |
| Range     | <b>NORMAL</b><br>Normal decimal encoding is used.  |
|           | <b>REVERSED</b><br>Reversed decimal encoding is used.  |
|           | For more information about encoding properties, see the WebSphere MQ <i>Using Java</i> document. |

**Floating point encoding:** If native encoding is not enabled, select the type of floating point encoding.

|           |  |
|-----------|--|
| Data type | Enum   |
| Units     | Not applicable   |
| Default   | IEEEENORMAL  |
| Range     | <b>IEEEENORMAL</b><br>IEEE normal floating point encoding is used.                               |
|           | <b>IEEEEVERSED</b><br>IEEE reversed floating point encoding is used.                             |
|           | <b>S390</b> S390 floating point encoding is used.  |
|           | For more information about encoding properties, see the WebSphere MQ <i>Using Java</i> document. |

**Native encoding:** Select this checkbox to indicate that the queue destination should use native encoding (appropriate encoding values for the Java platform).

|           |                |
|-----------|----------------|
| Data type | Enum           |
| Units     | Not applicable |
| Default   | Cleared        |

Range

**Cleared**

Native encoding is not used, so specify the properties below for integer, decimal, and floating point encoding.

**Selected**

Native encoding is used (to provide appropriate encoding values for the Java platform).

For more information about encoding properties, see the MQSeries *Using Java* document.

**Target client:** Whether the receiving application is JMS-compliant or is a traditional WebSphere MQ application

Data type

Enum

Units

Not applicable

Default

MQSeries

Range

**MQSeries**

The target is a non-JMS, traditional WebSphere MQ application.

**JMS**

The target is a JMS-compliant application.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### MQSeries topic destination settings for application clients

Use this panel to view or change the configuration properties of the selected topic destination for use with the MQSeries JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers** > *JMS provider instance*. Right-click **MQ Topic Destinations** and click **New**. The following fields appear on the **General** tab.

A topic destination is used to configure the properties of a JMS topic for the associated JMS provider. Connections to the topic are created by the associated topic connection factory. A topic for use with the MQSeries JMS provider has the following properties.

**Note:**

- The property values that you specify must match the values that you specified when configuring MQSeries JMS resources. For more information about configuring MQSeries JMS resources, see the MQSeries *Using Java* book.
- In MQSeries, names can have a maximum of 48 characters, with the exception of channels which have a maximum of 20 characters.

**Name:** The name by which the topic is known for administrative purposes.

Data type String

**Description:** A description of the topic, for administrative purposes within IBM WebSphere Application Server.

Data type String  
Default Null

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**Persistence:** Whether all messages sent to the destination are persistent, non-persistent, or have their persistence defined by the application

Data type Enum  
Units Not applicable  
Default APPLICATION\_DEFINED  
Range

**Application defined**

Messages on the destination have their persistence defined by the application that put them onto the queue.

**Queue defined**

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

**Persistent**

Messages on the destination are persistent.

**Non persistent**

Messages on the destination are not persistent.

**Priority:** Whether the message priority for this destination is defined by the application or the **Specified priority** property

Data type Enum  
Units Not applicable  
Default APPLICATION\_DEFINED  
Range

**Application defined**

The priority of messages on this destination is defined by the application that put them onto the destination.

**Queue defined**

[WebSphere MQ destination only] Messages on the destination have their persistence defined by the WebSphere MQ queue definition properties.

**Specified**

The priority of messages on this destination is defined by the **Specified priority** property. *If you select this option, you must define a priority on the **Specified priority** property.*

**Specified Priority:** If the **Priority** property is set to **Specified**, type here the message priority for this queue, in the range 0 (lowest) through 9 (highest)

If the **Priority** property is set to **Specified**, messages sent to this queue have the priority value specified by this property.

|           |  |
|-----------|--|
| Data type | Integer  |
| Units     | Message priority level                           |
| Default   | Null   |
| Range     | 0 (lowest priority) through 9 (highest priority) |

**Expiry:** Whether the expiry timeout for this queue is defined by the application or the **Specified expiry** property, or messages on the queue never expire (have an unlimited expiry timeout)

|           |                     |
|-----------|---------------------|
| Data type | Enum                |
| Units     | Not applicable      |
| Default   | APPLICATION_DEFINED |
| Range     |                     |

**Application defined**

The expiry timeout for messages on this queue is defined by the application that put them onto the queue.

**Specified**

The expiry timeout for messages on this queue is defined by the **Specified expiry** property. *If you select this option, you must define a timeout on the **Specified expiry** property.*

**Unlimited**

Messages on this queue have no expiry timeout, so those messages never expire.

**Specified Expiry:** If the **Expiry timeout** property is set to **Specified**, type here the number of milliseconds (greater than 0) after which messages on this queue expire

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | Milliseconds  |
| Default   | Null  |
| Range     | Greater than or equal to 0 <ul style="list-style-type: none"><li>• 0 indicates that messages never timeout</li><li>• Other values are an integer number of milliseconds</li></ul> |

**Base Topic Name:** The name of the topic to which messages are sent

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**CCSID:** The coded character set identifier for use with the WebSphere MQ queue manager.

This coded character set identifier (CCSID) must be one of the CCSIDs supported by WebSphere MQ.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Integer encoding:** If native encoding is not enabled, select whether integer encoding is normal or reversed.

|           |   |
|-----------|---|
| Data type | Enum  |
| Units     | Not applicable  |
| Default   | NORMAL  |
| Range     | <b>NORMAL</b><br>Normal integer encoding is used.     |
|           | <b>REVERSED</b><br>Reversed integer encoding is used. |

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

**Decimal encoding:** If native encoding is not enabled, select whether decimal encoding is normal or reversed.

|           |   |
|-----------|---|
| Data type | Enum  |
| Units     | Not applicable  |
| Default   | NORMAL  |
| Range     | <b>NORMAL</b><br>Normal decimal encoding is used.     |
|           | <b>REVERSED</b><br>Reversed decimal encoding is used. |

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

**Floating point encoding:** If native encoding is not enabled, select the type of floating point encoding.

|           |  |
|-----------|--|
| Data type | Enum   |
| Units     | Not applicable   |
| Default   | IEEEENORMAL  |
| Range     | <b>IEEEENORMAL</b><br>IEEE normal floating point encoding is used.   |
|           | <b>IEEEEVERSED</b><br>IEEE reversed floating point encoding is used. |
|           | <b>S390</b> S390 floating point encoding is used.                    |

For more information about encoding properties, see the WebSphere MQ *Using Java* document.

**Native encoding:** Select this checkbox to indicate that the queue destination should use native encoding (appropriate encoding values for the Java platform)..

|           |                |
|-----------|----------------|
| Data type | Enum           |
| Units     | Not applicable |
| Default   | Cleared        |

Range

**Cleared**

Native encoding is not used, so specify the properties above for integer, decimal, and floating point encoding.

**Selected**

Native encoding is used (to provide appropriate encoding values for the Java platform).

For more information about encoding properties, see the MQSeries *Using Java* document.

**BrokerDurSubQueue:** The name of the broker queue from which durable subscription messages are retrieved

The name of the broker queue from which durable subscription messages are retrieved. The subscriber specifies the name of the queue when it registers a subscription.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Default   | Null                          |
| Range     | 1 through 48 ASCII characters |

**BrokerCCDurSubQueue:** The name of the broker queue from which durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

The name of the broker queue from which durable subscription messages are retrieved for a ConnectionConsumer. This property applies only for use of the Web container.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Units     | En_US ASCII characters        |
| Default   | Null                          |
| Range     | 1 through 48 ASCII characters |

**Target Client:** Whether the receiving application is JMS-compliant or is a traditional MQSeries application

|           |                |
|-----------|----------------|
| Data type | Enum           |
| Units     | Not applicable |
| Default   | MQSeries       |
| Range     |                |

**MQSeries**

The target is a non-JMS, traditional MQSeries application.

**JMS**

The target is a JMS-compliant application.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property.

The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### Generic JMS connection factory settings for application clients

Use this panel to view or change the configuration properties of the selected JMS connection factory for use with the associated JMS provider. These configuration properties control how connections are created to the associated JMS destination.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers** > *new JMS Provider instance*. Right click **JMS Connection Factories** > click **New**. The following fields appear on the **General** tab.

A JMS connection factory is used to create connections to JMS destinations. The JMS connection factory is created by the associated JMS provider. A JMS connection factory for a generic JMS provider (other than the internal WebSphere JMS provider or the MQSeries JMS provider) has the following properties:

**Name:** The name by which this JMS connection factory is known for administrative purposes within IBM WebSphere Application Server. The name must be unique within the associated JMS provider.

Data type String

**Description:** A description of this connection factory for administrative purposes within IBM WebSphere Application Server.

Data type String  
Default Null

**JNDI Name:** The application client run-time uses this field to retrieve configuration information. The name must match the value of the **Name** field on the General tab in the Application Client Resource Reference section of the Application Assembly Tool.

**User:** The user ID used, with the **Password** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

The connection factory **User ID** and **Password** properties are used if the calling application does not provide a userid and password explicitly; for example, if the calling application uses the method `createQueueConnection()`. The JMS client flows the userid and password to the JMS server.

Data type String

**Password:** The password used, with the **User ID** property, for authentication if the calling application does not provide a userid and password explicitly.

If you specify a value for the **User ID** property, you must also specify a value for the **Password** property.

Data type String

Default

Null

**Re-Enter Password:** Confirms the password entered in the Password field.

**External JNDI Name:** The JNDI name that is used to bind the queue into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form *jms/Name*, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type

String

**Connection Type:** Whether this JMS destination is a queue (for point-to-point) or topic (for pub/sub).

Select one of the following options:

**Queue**

A JMS queue destination for point-to-point messaging.

**Topic** A JMS topic destination for pub/sub messaging.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

### Generic JMS destination settings for application clients

Use this panel to view or change the configuration properties of the selected JMS destination for use with the associated JMS provider.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **JMS Providers > new JMS Provider instance**. Right click **JMS Destinations > click New**. The following fields appear on the **General** tab.

A JMS destination is used to configure the properties of a JMS destination for the associated generic JMS provider. Connections to the JMS destination are created by the associated JMS connection factory. A JMS destination for use with a generic JMS provider (not the internal WebSphere JMS provider or MQSeries JMS provider) has the following properties.

**Name:** The name by which the queue is known for administrative purposes within IBM WebSphere Application Server.

Data type

String

**Description:** A description of the queue, for administrative purposes



**JNDI Name:** The JNDI name of the actual (physical) name of the JMS destination bound into JNDI.

**External JNDI Name:** The JNDI name that is used to bind the queue into the application server's name space.

As a convention, use the fully qualified JNDI name; for example, in the form `jms/Name`, where *Name* is the logical name of the resource.

This name is used to link the platform binding information. The binding associates the resources defined by the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

Data type String

**Destination Type:** Whether this JMS destination is a queue (for point-to-point) or topic (for pub/sub).

Select one of the following options:

**Queue**

A JMS queue destination for point-to-point messaging.

**Topic** A JMS topic destination for pub/sub messaging.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

**Example: Configuring JMS Provider, JMS Connection Factory and JMS Destination settings for application clients**

The purpose of this article is to help you to configure JMS Provider, JMS Connection Factory and JMS Destination settings.

- Required fields:
  - JMS Provider Properties page: name, and at least one protocol provider
  - JMS Connection Factory Properties page: name, jndiName, destination type
  - JMS Destination Properties page: name, jndiName, destination type
- Special cases:
  - The destination type must be QUEUE, or TOPIC.
- Example:

```
<resources.jms:JMSPProvider xmi:id="JMSPProvider_3" name="genericJMSPProvider:name"
description="genericJMSPProvider:description" externalInitialContextFactory=
"genericJMSPProvider:contextFactoryClass" externalProviderURL="genericJMSPProvider:providerUrl"
>
<classpath>genericJMSPProvider:classpath</classpath>
<factories xmi:type="resources.jms:GenericJMSPDestination" xmi:id="GenericJMSPDestination_1"
name="jmsDestination:name" jndiName="jmsDestination:jndiName" description=
"jmsDestination:description" externalJNDIName="jmsDestination:externalJndiName" type="QUEUE">
<propertySet xmi:id="J2EEResourcePropertySet_15">
<resourceProperties xmi:id="J2EEResourcePropertySet_17" name="jmsDestination:customName"
value="jmsDestination:customValue"/>
</propertySet>
</factories>
```

```

<factories xmi:type="resources.jms.GenericJMSConnectionFactory" xmi:id=
"GenericJMSConnectionFactory_1" name="jmsCF:name" jndiName="jmsCF:jndiName"
description="jmsCF:description" userID="jmsCF:user" password="{xor}NTIshB11MT4y0g==
" externalJNDIName="jmsCF:externalJndiName" type="QUEUE">
<propertySet xmi:id="J2EEResourcePropertySet_16">
<resourceProperties xmi:id="J2EEResourceProperty_18" name="jmsCF:customName"
value="jmsCF:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_17">
<resourceProperties xmi:id="J2EEResourceProperty_19" name="genericJMSProvider:customName"
value="genericJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

## Configuring new connection factories for application clients

### Before you begin

During this task, you create a new JMS connection factory configuration for your application client.

### Steps for this task

1. Click the JMS provider for which you want to create a connection factory in the tree. Do one of the following:
  - Configure a new JMS provider.
  - Click an existing JMS provider.
2. Expand the JMS provider to view its **JMS Connection Factories** folder.
3. Click the folder. Do one of the following:
  - Right-click the folder and click **New Factory**.
  - Click **Edit > New** on the menu bar.
4. Configure the JMS connection factory properties in the resulting property dialog.
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

## Configuring new Java Message Service destinations for application clients

### Before you begin

During this task, you create new Java Message Service (JMS) destination configuration for your application client.

### Steps for this task

1. Click the JMS provider in the tree for which you want to create a destination. Do one of the following:
  - Configure a new JMS provider.
  - Click an existing JMS provider.
2. Expand the JMS provider to view its **JMS Destinations** folder.
3. Click the folder. Do one of the following:
  - Right-click the folder and click **New Factory**.
  - Click **Edit > New** on the menu bar.
4. Configure the JMS destination properties in the resulting property dialog.

5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

## Example: Configuring MQ Queue and Topic connection factories and destination factories for application clients

The purpose of this article is to help you configure MQ Queue connection factory, MQ Topic connection factory, MQ Queue destination factory, and MQ Topic destination factory settings.

- Required fields:
  - MQ Queue Connection Factory Properties page: name, jndiName, transport type
  - MQ Topic Connection Factory Properties page: name, jndiName, broker version
  - MQ Queue Factory Properties page: name, jndiName, persistence, priority, expiry, baseQueueName, targetClient
  - MQ Topic Factory Properties page: name, jndiName, persistence, priority, expiry, baseQueueName, targetClient
- Special cases:
  - The transport type must be `CLIENT`, or `BINDINGS`.
  - The Broker Version must be `MA0C`, or `MQSI`.
  - The port must be a numerical value between -2417483648 and 2417483647.
  - The CCSID must be a numerical value between -2417483648 and 2417483647.
  - The persistence value must be `APPLICATION_DEFINED`, `QUEUE_DEFINED`, `PERSISTENT` or, `NONPERSISTENT`.
  - The priority must be `APPLICATION_DEFINED`, `QUEUE_DEFINED`, or `SPECIFIED`.
  - The expiry must be `APPLICATION_DEFINED`, `UNLIMITED`, or `SPECIFIED`.
  - The integer encoding must be `Normal`, or `Reversed`.
  - The decimal encoding must be `Normal`, or `Reversed`.
  - The floating encoding must be `IEEENormal`, `IEEEReversed`, `S390`.
  - The target client must be `JMS` or `MQ`.
  - On the MQ Queue Connection Factory Properties page, only set the `queueManager`, `host`, and `portWhen` (required) fields if the transport type is `CLIENT`.
  - On the MQ Topic Connection Factory Properties page, only set the `queueManager`, `host`, and `port` (required) fields if the transport type is `CLIENT`.
  - On the the MQ Topic Factory Properties, and the MQ Queue Factory Properties pages, only set the Integer encoding, decimal encoding, and floating point encoding (required) fields if you do not set `nativeEncoding`.
  - On the MQ Topic Factory Properties, and the MQ Queue Factory Properties pages, the specified priority entry field must be an integer between 0 and 9 if priority is set to `SPECIFIED`.
  - On the the MQ Topic Factory Properties, and the MQ Queue Factory Properties pages, the specified expiry entry field must be a value greater than 0 if expiry is set to `SPECIFIED`.
- Example:

```
<resources.jms:JMSProvider xmi:id="JMSProvider_1" name="MQ JMS Provider"
description="mqJMSProvider:description" externalInitialContextFactory=
"mqJMSProvider:contextFactoryClass" externalProviderURL="mqJMSProvider:providerUrl"
>
<classpath>mqJMSProvider:classpath</classpath>
```

```

<factories xmi:type="resources.jms.mqseries:MQQueueConnectionFactory"
xmi:id="MQQueueConnectionFactory_1" name="mqQCF:name" jndiName="mqQCF:jndiName"
description="mqQCF:description" userID="mqQCF:user" password="{xor}Mi40HB1lMT4y0g=="
" queueManager="mqQCF:queueManager" host="mqQCF:host" port="1" channel="mqQCF:channel"
transportType="CLIENT" clientID="mqQCF:clientID" CCSID="2">
<propertySet xmi:id="J2EEResourcePropertySet_3">
<resourceProperties xmi:id="J2EEResourceProperty_3" name="mqQCF:customName"
value="mqQCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQTopicConnectionFactory"
xmi:id="MQTopicConnectionFactory_1" name="mqTCF:name" jndiName="mqTCF:jndiName"
description="mqTCF:description" userID="mqTCF:user" password="{xor}Mi4LHB1lNTE7NhE+Mjo="
host="mqTCF:host" port="1" transportType="CLIENT" channel="mqTCF:channel" queueManager=
"mqTCF:queueManager" brokerControlQueue="mqTCF:brokerControlQueue" brokerQueueManager=
"mqTCF:brokerQueueManager" brokerPubQueue="mqTCF:brokerPubQueue" brokerSubQueue=
"mqTCF:brokerSubQueue" brokerCCSubQ="mqTCF:brokerCCSubQ" brokerVersion="MA0C"
clientID="mqTCF:clientID" CCSID="2">
<propertySet xmi:id="J2EEResourcePropertySet_4">
<resourceProperties xmi:id="J2EEResourceProperty_4" name="mqTCF:customName"
value="mqTCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQQueue" xmi:id="MQQueue_1" name="mqQ:name"
jndiName="mqQ:jndiName" description="mqQ:description" persistence="APPLICATION_DEFINED"
priority="SPECIFIED" specifiedPriority="1" expiry="SPECIFIED" specifiedExpiry="1"
baseQueueName="mqQ:baseQueueName" baseQueueManagerName="mqQ:baseQueueManagerName" CCSID="1"
integerEncoding="Normal" decimalEncoding="Normal" floatingPointEncoding="IEEENormal"
targetClient="JMS">
<propertySet xmi:id="J2EEResourcePropertySet_5">
<resourceProperties xmi:id="J2EEResourceProperty_5" name="mqQ:customName"
value="mqQ:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.mqseries:MQTopic" xmi:id="MQTopic_1" name="mqT:name"
jndiName="mqT:jndiName" description="mqT:description" persistence="APPLICATION_DEFINED"
priority="SPECIFIED" specifiedPriority="1" expiry="SPECIFIED" specifiedExpiry="2"
baseTopicName="mqT:baseTopicName" CCSID="3" integerEncoding="Normal" decimalEncoding="Normal"
floatingPointEncoding="IEEENormal" targetClient="JMS" brokerDurSubQueue="mqT:brokerDurSubQueue"
brokerCCDurSubQueue="mqT:brokerCCDurSubQueue">
<propertySet xmi:id="J2EEResourcePropertySet_6">
<resourceProperties xmi:id="J2EEResourceProperty_6" name="mqT:customName"
value="mqT:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_7">
<resourceProperties xmi:id="J2EEResourceProperty_7" name="mqJMSProvider:customName"
value="mqJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

## Example: Configuring WAS Queue and Topic connection factories and destination factories for application clients

The purpose of this article is to help you to configure WAS Queue connection factory, WAS Topic connection factory, WAS Queue destination factory, and WAS Topic destination factory settings.

- Required fields:
  - JMS Provider Properties page: name
  - WAS Queue Connection Factory Properties page: name, jndiName, node
  - WAS Topic Connection Factory Properties page: name, jndiName, node, port
  - WAS Queue Factory Properties page: name, jndiName, node, persistence, priority, expiry

- WAS Topic Factory Properties page: name, jndiName, topic name, persistence, priority, expiry
- Special cases:
  - The port must be QUEUED or DIRECT.
  - The CCSID must be a numerical value between -2417483648 and 2417483647.
  - The persistence value must be APPLICATION\_DEFINED, PERSISTENT, or NONPERSISTENT.
  - The priority must be APPLICATION\_DEFINED, or SPECIFIED.
  - The expiry must be APPLICATION\_DEFINED, UNLIMITED, or SPECIFIED.
  - On the WAS Topic Factory Properties, and the WAS Queue Factory Properties pages, the specified priority entry field must be an integer between 0 and 9 if priority is set to SPECIFIED .
  - On the WAS Topic Factory Properties, and the WAS Queue Factory Properties pages, the specified expiry entry field must be an value greater than 0 if expiry is set to SPECIFIED.
- Example:

```

<resources.jms:JMSProvider xmi:id="JMSProvider_2" name="WebSphere JMS Provider"
description="wasJMSProvider:description" externalInitialContextFactory=
"wasJMSProvider:contextfactoryclass" externalProviderURL="wasJMSProvider:providerUrl">
<classpath>wasJMSProvider:classpath</classpath>
<factories xmi:type="resources.jms.internalmessaging:WASQueueConnectionFactory"
xmi:id="WASQueueConnectionFactory_1" name="wasQCF:name" jndiName="wasQCF:jndiName"
description="wasQCF:description" userID="wasQCF:user" password="{xor}KD4sDhwZZSos0i0="
node="wasQCF:Node">
<propertySet xmi:id="J2EEResourcePropertySet_8">
<resourceProperties xmi:id="J2EEResourceProperty_8" name="wasQCF:customName"
value="wasQCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASTopicConnectionFactory"
xmi:id="WASTopicConnectionFactory_1" name="wasTCF:name" jndiName="wasTCF:jndiName"
description="wasTCF:description" userID="wasTCF:user" password="{xor}KD4sCxwZZTE+Mjo="
node="wasTCF:node" port="QUEUED" clientID="wasTCF:clientId">
<propertySet xmi:id="J2EEResourcePropertySet_9">
<resourceProperties xmi:id="J2EEResourceProperty_9" name="wasTCF:customName"
value="wasTCF:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASQueue" xmi:id="WASQueue_1"
name="wasQ:name" jndiName="wasQ:jndiName" description="wasQ:description" node="wasQ:node"
persistence="APPLICATION_DEFINED" priority="SPECIFIED" specifiedPriority="1" expiry="SPECIFIED"
specifiedExpiry="1">
<propertySet xmi:id="J2EEResourcePropertySet_10">
<resourceProperties xmi:id="J2EEResourceProperty_10" name="wasQ:customName"
value="wasQ:customValue"/>
</propertySet>
</factories>
<factories xmi:type="resources.jms.internalmessaging:WASTopic" xmi:id="WASTopic_1"
name="wasT:name" jndiName="wasT:jndiName" description="wasT:description" topic="wasT:topicName"
persistence="APPLICATION_DEFINED" priority="SPECIFIED" specifiedPriority="1" expiry="SPECIFIED"
specifiedExpiry="1">
<propertySet xmi:id="J2EEResourcePropertySet_11">
<resourceProperties xmi:id="J2EEResourceProperty_11" name="wasT:customName"
value="wasT:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_12">
<resourceProperties xmi:id="J2EEResourceProperty_12" name="wasJMSProvider:customName"
value="wasJMSProvider:customValue"/>
</propertySet>
</resources.jms:JMSProvider>

```

## Configuring new resource environment providers for application clients

### Before you begin

During this task, you create new resource environment provider configurations for your application client.

To configure a new resource environment provider, perform the following steps:

### Steps for this task

1. Start the tool and open the EAR file for which you want to configure the new JMS provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file in which you want to configure the new JMS provider.
3. Expand the JAR file to view its contents.
4. Click the folder called **Resource Environment Providers**. Do one of the following:
  - Right-click the folder and click **New Provider**.
  - Click **Edit > New** on the menu bar.
5. Configure the JMS provider properties in the resulting property dialog.
6. Click **OK** when finished.
7. Click **File > Save** on the menu bar to save your changes.

### **Resource environment provider settings for application clients**

Use this page to specify resource environment entry properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file. Right-click **Resource Environment Providers >** and click **New**. The following fields appear on the **General** tab:

**Name:** Specifies the administrative name for the resource environment provider.

**Description:** Specifies a description of the resource environment provider for your administrative records.

**Class Path:** Specifies the path to the JAR file that contains the implementation classes for the resource environment provider.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

## Configuring new resource environment entries for application clients

### Before you begin

During this task, you create new resource environment entries for your client application.

### Steps for this task

1. Click the resource environment provider in the tree, for which you want to create a resource environment entry. Do one of the following:
  - Configure a new resource environment provider.
  - Click an existing resource environment provider.
2. Expand the resource environment provider to view the **resource environment entries** folder.
3. Click the folder. Do one of the following:
  - Right-click the folder and select **New Factory**.
  - Click **Edit > New** on the menu bar.
4. Configure the data source properties in the resulting property dialog.
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

### **Resource environment entry settings for application clients**

Use this page to specify resource environment entry properties.

To view this Application Client Resource Configuration Tool (ACRCT) page, click **File > Open**. After you browse for an EAR file, click **Open**. Expand the selected JAR file > **Resource Environment Providers** > *resource environment instance*. Right-click **Resource environment entry** > and click **New**. The following fields appear on the **General** tab:

**Name:** Specifies the administrative name for the resource environment entry.

**Description:** Specifies a description of the URL for your administrative records.

**JNDI Name:** Specifies the Java Naming and Directory Interface (JNDI) name for the resource, including any naming subcontexts.

Use this name to link to the binding information of the platform. The binding associates the resources defined in the deployment descriptor of the module to the actual (physical) resources bound into JNDI by the platform.

**Custom Properties:** Specifies name-value pairs for setting additional properties on the object that is created at runtime for this resource.

You must enter a name that is a public property on the object and a value that can be converted from a string to the type required by the set method of the property. The acceptable properties and values depend on the object that is created. Refer to the object documentation for a list of valid properties and values.

---

## **Managing application clients**

### Before you begin

Perform the following tasks after deploying application clients.

**Note:** This task only applies to J2EE application clients.

### Steps for this task

1. Update data source and data source provider configurations.
2. Update URLs and URL provider configurations.

3. Update mail session configurations.
4. Update JMS provider, connection factories, and destination configurations.
5. Update MQ JMS provider, MQ connection factories, and MQ destination configurations.
6. Update Resource Environment Entry and Resource Environment Provider configurations.
7. (Optional) Remove application client resources.

## Updating data source and data source provider configurations with the Application Client Resource Configuration Tool

### Before you begin

During this task, you update the configuration of an existing data source or data source provider.

### Steps for this task

1. Start the tool and open the EAR file containing the data source or data source provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the data source or data source provider to update.
3. Expand the JAR file to view its contents until you locate the particular data source or data source provider to update. Do one of the following:
  - Right-click the object and click **Properties**.
  - Click **Edit > Properties** on the menu bar.
4. Update the properties in the resulting property dialog. For detailed field help, go to:
  - Data source provider properties
  - Data source properties
5. Click **OK** when finished.
6. Click **File > Save** on the menu bar to save your changes.

## Updating URLs and URL provider configurations for application clients

### Steps for this task

1. Start the tool and open the EAR file containing the URL or URL provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the URL or URL provider to update.
3. Expand the JAR file to view its contents.
4. Keep expanding the JAR file contents until you locate the particular URL or URL provider to update. Do one of the following:
  - a. Right-click the object and click **Properties**
  - b. Click **Edit > Properties** on the menu bar.
5. Update the properties in the resulting property dialog.
6. Click **OK** when finished.
7. Click **File > Save** to save your changes on the menu bar.



## Updating mail session configurations for application clients

### Before you begin

During this task, you update the configuration of an existing JavaMail session.

### **Note:**

You cannot update the name of the default JavaMail provider. Also, you cannot delete the default JavaMail provider from the tree.

### Steps for this task

1. Start the tool and open the EAR file containing the JavaMail session.  
The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the JavaMail session to update.
3. Expand the JAR file to view its contents.
4. Keep expanding the JAR file contents until you locate the particular JavaMail session to update. Do one of the following:
  - a. Right-click the object and click **Properties**
  - b. Click **Edit > Properties** from the menu bar.
5. Update the properties in the resulting property dialog.
6. Click **OK** when finished.
7. Select **File > Save** from the menu bar to save your changes.

## Updating Java Message Service provider, connection factories, and destination configurations for application clients

### Before you begin

During this task, you update the configuration of an existing Java Message Service (JMS) provider, connection factory, or destination.

### Steps for this task

1. Start the tool and open the EAR file containing the JMS provider, connection factory, or destination. The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the JMS provider, connection factory, or destination to update.
3. Expand the JAR file to view its contents until you locate the particular JMS provider, connection factory, or destination to update. When you find it, do one of the following:
  - Right-click the object and click **Properties**.
  - Click **Edit > Properties** on the menu bar.
4. Update the properties in the resulting property dialog. For detailed field help, see:
  - JMS provider properties
  - WAS Queue connection factory properties
  - WAS Topic connection factory properties
  - WAS Queue destination properties
  - WAS Topic destination properties
5. Click **OK**.

6. Click **File > Save** to save your changes.

## Updating MQ Java Message Service provider, MQ connection factories, and MQ destination configurations for application clients

### Before you begin

During this task, you will update the configuration of an existing MQ JMS provider, MQ connection factory, or MQ destination.

### Steps for this task

1. Start the tool and open the EAR file containing the MQ JMS provider, MQ connection factory, or MQ destination. The EAR file contents will be displayed in a tree view.
2. Select from the tree the JAR file containing the MQ JMS provider, MQ connection factory, or MQ destination to update.
3. Expand the JAR file to view its contents until you locate the particular MQ JMS provider, MQ connection factory, or MQ destination that you want to update. Do one of the following:
  - Right-click the object and click **Properties**.
  - Click **Edit > Properties** on the menu bar.
4. Update the properties in the resulting property dialog. For detailed field help, see:
  - JMS provider properties
  - MQ Queue connection factory properties
  - MQ Topic connection factory properties
  - MQ Queue destination properties
  - MQ Topic destination properties
5. Click **OK**.
6. Click **File > Save** to save your changes.

## Updating Resource Environment Entry and Resource Environment Provider configurations for application clients

### Before you begin

During this task, you update the configuration of an existing Resource Environment Entry or Resource Environment Provider.

### Steps for this task

1. Start the tool and open the EAR file containing the Resource Environment Entry or Resource Environment Provider. The EAR file contents display in a tree view.
2. Select from the tree the JAR file containing the Resource Environment Entry or Resource Environment provider to update.
3. Expand the JAR file to view its contents until you locate the Resource Environment Entry or Resource Environment Provider to update. Do one of the following:
  - Right-click the object and click **Properties**.
  - Click **Edit > Properties** on the menu bar.

4. Update the properties in the resulting property dialog. For detailed field help, see:
  - Resource environment provider properties
  - Resource environment entry properties
5. Click **OK** when you finish.
6. Click **File > Save** on the menu bar to save your changes.

### Example: Configuring Resource Environment settings

The purpose of this article is to help you configure Resource Environment settings.

- Required fields:
  - Resource Environment Provider page: name
  - Resource Environment Entry page: name, jndiName
- Example:

```
<resources.env:ResourceEnvironmentProvider xmi:id="ResourceEnvironmentProvider_1"
name="resourceEnvProvider:name" description="resourceEnvProvider:description">
<classpath>resourceEnvProvider:classpath</classpath>
<factories xmi:type="resources.env:ResourceEnvEntry" xmi:id="ResourceEnvEntry_1"
name="resourceEnvEntry:name" jndiName="resourceEnvEntry:jndiName" description="resourceEnvEntry:
description">
<propertySet xmi:id="J2EEResourcePropertySet_20">
<resourceProperties xmi:id="J2EEResourceProperty_22" name="resourceEnvEntry:customName"
value="resourceEnvEntry:customValue"/>
</propertySet>
</factories>
<propertySet xmi:id="J2EEResourcePropertySet_21">
<resourceProperties xmi:id="J2EEResourceProperty_23" name="resourceEnvProvider:customName"
value="resourceEnvProvider:customValue"/>
</propertySet>
</resources.env:ResourceEnvironmentProvider>
```

### Example: Configuring Resource Environment custom settings for application clients

The purpose of this article is to help you configure Resource Environment custom settings.

- The custom page applies to every resource type. You can specify as many custom names and values as you need.
- Example:

```
<propertySet xmi:id="J2EEResourcePropertySet_20">
<resourceProperties xmi:id="J2EEResourceProperty_22"
name="resourceEnvEntry:customName"
value="resourceEnvEntry:customValue"/>
</propertySet>
```

## Removing application client resources

Before you begin

**Note:** This task only applies to J2EE application clients.

Steps for this task

1. Start the Application Client Resource Configuration Tool (ACRCT) and open the EAR file from which you want to remove an object. The EAR file contents display in a tree view.
 

If you already have an EAR file open, and have made some changes, click **File > Save** to save your work before proceeding to delete an object.
2. Locate the object that you want to remove in the tree.

3. Right-click the object, then click **Delete**.
4. Click **File > Save**.

#### What to do next

The option to delete an item does not offer a confirmation dialog. As a safeguard, consider saving your work right before you begin this task. If you change your mind after removing an item, you can close the EAR file without saving your changes, canceling your deletion. Remember to close the EAR file immediately after the deletion, or you also lose any unsaved work that you performed since the deletion.

---

## Running application clients

The J2EE specification requires support for a client container that runs stand-alone Java applications (known as J2EE application clients) and provides J2EE services to the applications. J2EE services include naming, security, and resource connections.

You are ready to run your application client using this tool after you have:

1. Written the application client program.
2. Assembled and installed an application module (.ear file) in the application server run time.
3. Deployed the application using the Application Client Resource Configuration Tool (ACRCT).

**Note:** This task only applies to J2EE application clients.

#### Steps for this task

1. Open a command window and invoke the following script to launch J2EE application clients using the launchClient shell:

```
install_root/bin/launchClient.bat
```

The launchClient batch command starts the application client run time, which:

- Initializes the client run time.
- Loads the class that you designated as the main class with the Application Assembly Tool (AAT).
- Runs the main method of the application client program.

When your program terminates, the application client run time cleans up the environment and the Java Virtual Machine code ends.

2. Pass parameters to the **launchClient** command. You can pass parameters to your application client program as well. The launchClient command allows you to do both.

The launchClient command requires that the first parameter is either:

- An EAR file specifying the application client to launch.
- A request for launchClient usage information.

All other parameters intended for the **launchClient** command must begin with the -CC prefix.

Parameters that are not EAR files, or usage requests, or that do not begin with the -CC prefix, are ignored by the application client run time, and are passed directly to the application client program.

The **launchClient** command retrieves parameters from three places:

- a. The command line
- b. A properties file
- c. System properties

The parameters are resolved in the order listed above, with command line values having the highest priority and system properties the lowest. This prioritization allows you to set and override default values.

### 3. Specify the server name.

By default, the **launchClient** command uses the environment variable `COMPUTERNAME` for the `BootstrapHost` property value. This setting is effective for testing your application client when it is installed on the same computer as the server. However, in other cases override this value with the name of your server.

You can override the `BootstrapHost` value by invoking `launchClient` with the following parameters:

```
launchClient myapp.ear -CCBootstrapHost=abc.midwest.mycompany.com
```

You can also override the default by specifying the value in a properties file and passing the file name to the `launchClient` shell.

**Note:** Security is controlled by the server. You do not need to configure security on the client because the client assumes that security is enabled. If security is not enabled, the server ignores the security request, and the application client works as expected.

### Usage scenario

You can store `launchClient` values in a properties file, a good method for distributing default values. You can then override one or more values on the command line. The format of the file is one `launchClient -CC` parameter per line without the `-CC` prefix. For example:

```
verbose=true  
classpath=c:\mydir\util.jar;c:\mydir\harness.jar;c:\production\G19\global.jar  
BootstrapHost=abc.westcoast.mycompany.com tracefile=c:\WebSphere\mylog.txt
```

## launchClient tool

### Syntax

This section describes the command line syntax for the Java TM2 Platform, WebSphere Application Server Enterprise (J2EE) `launchClient` tool.

The command line invocation syntax for the `launchClient` tool follows:

```
launchClient [<userapp.ear> |-help|-?] [-CCname=  
value] [app args]
```

where *userapp.ear* is the path and the name of the EAR file that contains the application client, *name* is the name of the parameter, *value* is the value to which the parameter ID is set, and *app args* are arguments that pass to the application client.

To print the usage information, the first parameter must be a path and a name to an EAR file, `-help`, or `-?`. All other parameters are optional and can appear in any

order. The application client run time ignores any optional parameters that do not begin with a -CC prefix, and passes them to the application client.

## Parameters

Supported arguments include:

### **-CCsoapConnectorPort**

The soap connector port. If you do not specify this argument, the WebSphere Application Server default value is used.

### **-CCverbose**

This option displays additional information messages. The default is false.

### **-CCclasspath**

A class path value. When you launch an application, the system class path is not used. If you want to access classes that are not in the EAR file or part of the resource class paths, specify the appropriate class path here. Multiple paths can be concatenated.

### **-CCjar**

The name of the client JAR file that resides within the EAR file for the application you wish to launch. Use this argument when you have multiple client JAR files in the EAR file.

### **-CCadminConnectorHost**

Specifies the host name of the server from which configuration information is retrieved. The default is the value of the -CCBootstrapHost parameter or the value of the local host if the -CCBootstrapHost parameter is not specified.

### **-CCadminConnectorPort**

Indicates the port number that the administrative client function should use. The default value is 8880 for SOAP connections and 2809 for RMI connections.

### **-CCadminConnectorType**

Specifies how the administrative client should connect to the server. Specify RMI to use the RMI connection type or specify SOAP to use the SOAP connection type. The default value is SOAP.

### **-CCadminConnectorUser**

Administrative clients use this user name when a server requires authentication. If the connection type is SOAP, and security is enabled on the server, this parameter is required. The SOAP connector does not prompt for authentication.

### **-CCadminConnectorPassword**

The password for the user name that the -CCadminConnectorUser parameter specifies.

### **-CCaltDD**

The name of an alternate deployment descriptor. This parameter is used with the -CCjar parameter to specify the deployment descriptor to use. Use this argument when a client jar file is configured with more than one deployment descriptor. Set the value to null to use the client JAR file standard deployment descriptor.

### **-CCBootstrapHost**

The name of the host server you want to connect to initially. The format is: *<i>your.server.ofchoice.com</i>*

**-CCBootstrapPort**

The server port number. If you do not specify this argument, the WebSphere Application Server default value is used.

**-CCproviderURL**

Provides bootstrap server information that the initial context factory can use to obtain an initial context. WebSphere Application Server initial context factory can use either a CORBA object URL or an IIOP URL. CORBA object URLs are more flexible than IIOP URLs and are the recommended URL format to use. This value can contain more than one bootstrap server address. This feature can be used when attempting to obtain an initial context from a server cluster. You can specify bootstrap server addresses, for all servers in the cluster, in the URL. The operation will succeed if at least one of the servers is running, eliminating a single point of failure. The address list does not process in a particular order. For naming operations, this value overrides the `-CCBootstrapHost` and `-CCBootstrapPort` parameters. An example of a CORBA object URL specifying multiple systems follows:

```
-CCproviderURL=  
  corbaloc:iiop:myserver.mycompany.com:9810,:mybackupserver.mycompany  
  .com:2809
```

This value is mapped to the `java.naming.provider.url` system property.

**-CCinitonly**

Use this option to initialize application client run time for ActiveX application clients without launching the client application. The default is false.

**-CCtrace**

Use this option to obtain debug trace information. You might need this information when reporting a problem to IBM Service. The default is false.

**-CCtracefile**

The name of the file to write trace information. The default is to output to the console.

**-CCpropfile**

Name of a properties file that contains launchClient properties. Specify the properties without the `-CC` prefix in the file. For example: `verbose=true`.

**-CCsecurityManager**

Enables and runs the WebSphere Application Server with a security manager. The default is disable.

**-CCsecurityMgrClass**

The fully qualified name of a class that implements a security manager. Only use this argument if the `-CCsecurityManager` parameter is set to enable. The default is `java.lang.SecurityManager`.

**-CCsecurityMgrPolicy**

The name of a security manager policy file. Only use this argument if the `-CCsecurityManager` parameter is set to enable. When you enable this parameter, the `java.security.policy` system property is set. The default is `<install_root>/properties/client.policy`.

**-CCD** Use this option to have the WebSphere Application Server set the specified system property during initialization. Do not use the `=` character after the `-CCD`. For example: `-CCDcom.ibm.test.property=testvalue`. You can specify

multiple -CCD parameters. The general format of this parameter is -CCD<property key>=<property value>.

#### **-CCexitVM**

Use this option to have the WebSphere Application Server call `System.exit()` after the client application completes. The default is `false`.

#### **-CCdumpJavaNameSpace**

Prints out the Java portion of the WebSphere Application Server Java Naming and Directory Interface (JNDI) name space. The `true` value uses the short format which prints out the binding name and the type of the object bound at that location. The `long` value uses the long format which prints out the binding name, bound object type, local object, type, and string representation of the local object, for example: IORs, and string values. The default value is `false`.

### **Examples**

The following examples demonstrate correct syntax.

#### **On the Windows operating system:**

```
launchClient c:\earfiles\myapp.ear -CCBootstrapHost=myWASServer  
-CCverbose=true app_parm1 app_parm2
```

#### **On the UNIX operating system:**

```
./launchClient.sh /usr/earfiles/myapp.ear  
-CCBootstrapHost=myWASServer -CCverbose=true app_parm1 app_parm2
```

### **Specifying the directory for an expanded EAR file**

#### Before you begin

Each time `launchClient` is called, it extracts the EAR file to a random directory name in the temporary directory on your hard drive. Then it sets up the thread `ClassLoader` to use the extracted EAR file directory and JAR files included in the `Manifest.mf` client JAR file. In a normal J2EE Java client, these files are automatically cleaned up after the application exits. This cleanup occurs when the client container shutdown hook is called. To avoid extracting the EAR file (and removing the temporary directory) each time `launchClient` is called, complete the following steps:

#### Steps for this task

1. Specify a directory to extract the EAR file by setting the `com.ibm.websphere.client.applicationclient.archivedir` Java system property. If the directory does not exist or is empty, the EAR file is extracted normally. If the EAR file was previously extracted, the `launchClient` tool reuses the directory.
2. Delete the directory before running the `launchClient` tool again, if you need to update your EAR file.

When you call the `launchClient` command, it extracts the new EAR file to the directory. If you do not delete the directory or change the system property value to point to a different directory, `launchClient` reuses the currently extracted EAR file, and does not use your changed EAR file.

**Note:** When specifying the `com.ibm.websphere.client.applicationclient.archivedir` property, make sure that the directory you specify is unique for each EAR file you use. For example, do not point `MyEar1.ear` and `MyEar2.ear` files to the same directory.



## Example: Using a Java 2 security manager with a J2EE application client

The `launchClient` command provides several parameters to control the use of a Java 2 security manager. By default the `launchClient` command does not enable nor run with a Java 2 security manager. To enable the Java 2 security manager, add the following parameter to your `launchClient` command:

```
-CCsecurityManager=enable
```

For example:

```
launchClient myear.ear -CCsecurityManager=enable
```

- When the security manager is enabled, the `launchClient` uses by default the `java.lang.SecurityManager` class and the `<WAS_HOME>/properties/client.policy` policy file. This policy file is configured to provide the standard permissions as described in the J2EE specification for J2EE application clients and applets. If your application receives a `java.security.AccessControlException`, you must add additional permissions to the `client.policy` file. For more information on adding additional permissions, see configuring `client.policy` files and `AccessControlException`.
- You can override the default security manager class by specifying the `-CCsecurityMgrClass` parameter and the default policy file using the `-CCsecurityMgrPolicy` parameter. For more information, see `launchClient` tool.
- If you invoke Java to start the `launchClient` class, it is recommended that you do not use the `-Djava.security.manager` parameter to enable the Java 2 security manager. Using this parameter causes the Java 2 security manager to be enabled prior to initialization of the J2EE application client runtime. The necessary permissions are not granted and your application may receive `java.security.AccessControlExceptions`.
- When the J2EE application client runtime is initialized, the Enterprise Archive (EAR) file that you specified is extracted to a random subdirectory in your users temporary directory location.

**Note:** If the EAR file is a set of directories and subdirectories, then it is used in place and not expanded. The J2EE application client runtime sets the `com.ibm.websphere.client.applicationclient.archivedir` system property to the directory location of the EAR file. The `client.policy` file uses this system property to inform the security manager of the location of your application client codebase and to assign the configured permissions to that codebase. This occurs when the security manager is enabled. If the security manager is enabled at the time Java is started, then this system property is not set, the codebase is unknown, and the permissions can not be granted.

- It is recommended that you enable the security manager with the J2EE application client runtime. Use the following parameter:  
`-CCsecurityManager=enable`.

## Example: Enabling Java 2 security prior to J2EE application client runtime initialization

To enable the Java 2 security prior to the J2EE application client runtime initialization, set the `com.ibm.websphere.client.applicationclient.archivedir` system property. Perform the following steps:

1. Set the system property to directory where the Enterprise Archive (EAR) should be expanded to, for example:  
`-Dcom.ibm.websphere.client.applicationclient.archivedir=c:\myear1` (windows)  
`-Dcom.ibm.websphere.client.applicationclient.archivedir=/usr/mrear1` (Unix)

2. Set the `java.security.policy` system property to use the `<WAS_HOME>/properties/client.policy` file, for example:
  - `Djava.security.policy=%WAS_HOME%\properties\client.policy` (Windows)
  - `Djava.security.policy=$WAS_HOME/properties/client.policy` (Unix)
- Setting the `com.ibm.websphere.client.applicationclient.archivedir` has the following effects:
  - If the directory does not exist or it is empty, the EAR file is extracted to that directory.
  - The EAR file is reused if it was previously extracted. This occurs even if the EAR file specified on the command line is different.
  - The security manager will grant the permissions from the `client.policy` file to that directory and all its subdirectories.
- There are two types of EAR files. The first type of EAR file is a single file that contains all the enterprise application files. The second type is a set of directories and subdirectories. The following only applies if you are using the single file form:
  - If you need to update your EAR file, delete the directory first.
  - The new EAR file will be extracted to the directory the next time you run. If you do not delete the directory or change the system property value to point to a different temporary directory, the currently extracted EAR file will be reused, and your changed EAR file will not be used.
- When specifying the `com.ibm.websphere.client.applicationclient.archivedir` property, ensure that the directory you specify is unique for each EAR file that you use. For example, do not point `MyEar1.ear` and `MyEar2.ear` files to the same directory. You also must create all directories up to, but not including, the last directory. For example, if you set the following:
 

```
com.ibm.websphere.client.applicationclient.archivedir=/usr/myears/myear1
```

then `<i>usr</i>` and `<i>myears</i>` must exist, but `<i>myear1</i>` does not have to exist prior to running the `launchClient` class.

---

## Application client troubleshooting tips

This section provides some debugging tips for resolving common J2EE application client problems. To use this troubleshooting guide, review the trace entries for one of the J2EE application client exceptions, and then locate the exception in the guide. Some of the errors in the guide are samples; the actual error you receive can be slightly different than what is shown here. Also, it can be useful to rerun the **launchClient** command specifying the `-CCverbose=true` option. This option provides additional information when the J2EE application client run time is initializing

### Error: `java.lang.NoClassDefFoundError`

|                 |  |
|-----------------|--|
| Explanation     | This exception is thrown when Java code cannot load the specified class.   |
| Possible causes | <ul style="list-style-type: none"> <li>• Invalid or non-existent class</li> <li>• Classpath problem</li> <li>• Manifest problem</li> </ul> |

## Recommended response

Check to determine if the specified class exists in a JAR file within your EAR file. If it does, make sure the path for the class is correct. For example, if you get the exception:

```
java.lang.NoClassDefFoundError:  
WebSphereSamples.HelloEJB.HelloHome
```

ensure the class HelloHome exists in one of the JAR files in your EAR file. If it exists, ensure the path for the class is WebSphereSamples.HelloEJB.

If both the class and path are correct, then it is a classpath issue. Most likely, you do not have the failing class JAR file specified in the client JAR file manifest. To verify this situation, perform the following steps:

1. Open your EAR file with the Application Assembly Tool and click on the Application Client.
2. Add the names of the other JAR files in the EAR file to the Classpath field.

This exception is generally caused by a missing EJB module name from the Classpath field.

If you have multiple JAR files to enter in the Classpath field, be sure to separate the JAR names with spaces.

If you still have the problem, you have a situation where a class is loaded from the file system instead of the EAR file. This is a very difficult situation to debug because the offending class is not the one specified in the exception. Instead, another class is loaded from the file system before the one specified in the exception. To correct this problem, review the classpaths specified with the `-CClasspath` option and the classpaths configured with the Application Client Resource Configuration Tool. Look for classes that also exist in the EAR file. You must resolve the situation where one of the classes is found on the file system instead of in the `.ear` file. Remove entries from the classpaths, or include the `.jar` files and classes in the `.ear` file instead of referencing them from the file system.

If you use the `-CClasspath` parameter or resource classpaths in the Application Client Resource Configuration Tool, and you have configured multiple JAR files or classes, verify they are separated with the correct character for your operating system. Unlike the classpath field in the Application Assembly Tool, these classpath fields use platform-specific separator characters, usually a colon (on UNIX platforms) or a semi-colon (on Windows systems).

**Note:** The system classpath is not used by the Application Client run time if you use the `launchClient` batch or shell files. In this case, the system classpath would not cause this problem. However, if you load the `launchClient` class directly, you do have to search through the system classpath as well.

**Error: com.ibm.websphere.naming.CannotInstantiateObjectException: Exception occurred while attempting to get an instance of the object for the specified reference object. [Root exception is javax.naming.NameNotFoundException: xxxxxxxxxx]**

|                      |  |
|----------------------|--|
| Explanation          | This exception occurs when you perform a lookup on an object that is not installed on the host server. Your program can look up the name in the local client Java Naming and Directory Interface (JNDI) name space, but received a <code>NameNotFoundException</code> exception because it is not located on the host server. One typical example is looking up an enterprise bean that is not installed on the host server that you access. This exception might also occur if the JNDI name you configured in your Application Client module does not match the actual JNDI name of the resource on the host server.   |
| Possible causes      | <ul style="list-style-type: none"> <li>• Incorrect host server invoked</li> <li>• Resource is not defined</li> <li>• Resource is not installed</li> <li>• Application server is not started</li> <li>• Invalid JNDI configuration</li> </ul>   |
| Recommended response | If you are accessing the wrong host server, run the <b>launchClient</b> command again with the <code>-CCBootstrapHost</code> parameter specifying the correct host server name. If you are accessing the correct host server, use the WebSphere <b>dumpnamespace</b> command line tool to see a listing of the host server JNDI name space. If you do not see the failing object name, the resource is either not installed on the host server or the appropriate application server is not started. If you determine the resource is already installed and started, your JNDI name in your client application does not match the global JNDI name on the host server. Use the Application Assembly Tool to compare the JNDI bindings value of the failing object name in the client application to the JNDI bindings value of the object in the host server application. They must match. |

**Error: javax.naming.ServiceUnavailableException: A communication failure occurred while attempting to obtain an initial context using the provider url: "iiop://[invalidhostname]". Make sure that the host and port information is correct and that the server identified by the provider URL is a running name server. If no port number is specified, the default port number 2809 is used. Other possible causes include the network environment or workstation network configuration. Root exception is org.omg.CORBA.INTERNAL: JORB0050E: In Profile.getIPAddress(), InetAddress.getByName[invalidhostname] threw an UnknownHostException. minor code: 4942F5B6 completed: Maybe**

|                      |  |
|----------------------|--|
| Explanation          | This exception occurs when you specify an invalid host server name.  |
| Possible causes      | <ul style="list-style-type: none"> <li>• Incorrect host server invoked</li> <li>• Invalid host server name</li> </ul>                        |
| Recommended response | Run the <b>launchClient</b> command again and specify the correct name of your host server with the <code>-CCBootstrapHost</code> parameter. |

**Error: javax.naming.CommunicationException: Could not obtain an initial context due to a communication failure. Since no provider URL was specified, either the bootstrap host and port of an existing ORB was used, or a new ORB instance was created and initialized with the default bootstrap host of "localhost" and the default bootstrap port of 2809. Make sure the ORB bootstrap**

**host and port resolve to a running name server. Root exception is org.omg.CORBA.COMM\_FAILURE: WRITE\_ERROR\_SEND\_1 minor code: 49421050 completed: No**

|                      |   |
|----------------------|---|
| Explanation          | This exception occurs when you run the <b>launchClient</b> command to a host server that does not have the Application Server started. You also receive this exception when you specify an invalid host server name. This situation might occur if you do not specify a host server name when you run <b>launchClient</b> . The default behavior is for <b>launchClient</b> to run to <b>localhost</b> , because WebSphere Application Server does not know the name of your host server. This default behavior only works when you are running the client on the same computer with WebSphere Application Server is installed. |
| Possible causes      | <ul style="list-style-type: none"><li>• Incorrect host server invoked</li><li>• Invalid host server name</li><li>• Invalid reference to <b>localhost</b></li><li>• Application server is not started</li><li>• Invalid bootstrap port</li></ul>   |
| Recommended response | If you are not running to the correct host server, run the <b>launchClient</b> command again and specify the name of your host server with the <b>-CCBootstrapHost</b> parameter. Otherwise, start the Application Server on the host server and run the <b>launchClient</b> command again.   |

**Error: javax.naming.NameNotFoundException: Name comp/env/ejb not found in context "java:"**

|                      |   |
|----------------------|---|
| Explanation          | This exception is thrown when the Java code cannot locate the specified name in the local JNDI name space.  |
| Possible causes      | <ul style="list-style-type: none"><li>• No binding information for the specified name</li><li>• Binding information for the specified name is incorrect</li><li>• Wrong class loader was used to load one of the program classes</li><li>• A resource reference does not include any client configuration information</li><li>• A client container on the deployment manager is trying to use enterprise extensions (not supported)</li></ul>   |
| Recommended response | Open the EAR file with the Application Assembly Tool and check the bindings for the failing name. Ensure this information is correct. If you are using Resource References, open the EAR file with the Application Client Resource Configuration Tool, and make sure the Resource Reference has client configuration information and the name of the Resource Reference exactly matches the JNDI name of the client configuration. If it is correct, you might have a class loader issue. |

**Error: java.lang.ClassCastException: Unable to load class: org.omg.stub.WebSphereSamples.HelloEJB\_HelloHome\_Stub at com.ibm.rmi.javax.rmi.PortableRemoteObject.narrow(portableRemoteObject.java:269)**

|                      |  |
|----------------------|--|
| Explanation          | This exception occurs when the application program attempts to narrow to the EJB home class and the class loaders cannot find the EJB client side bindings.  |
| Possible causes      | <ul style="list-style-type: none"> <li>• The files, *_Stub.class and _Tie.class, are not in the EJB .jar file</li> <li>• Class loader could not find the classes</li> </ul>  |
| Recommended response | Look at the EJB .jar file located in the .ear file and verify the class contains the EJB client side bindings. These are class files whose names end in _Stub and _Tie. If these files are not present, then use the Application Assembly Tool to generate the binding classes. For more information, see article Generating deployment code for modules. If the binding classes are in the EJB .jar file, then you might have a class loader issue. |

**Error: WSCL0210E: The Enterprise archive file [EAR file name] could not be found. com.ibm.websphere.client.applicationclient.ClientContainerException: com.ibm.etools.archive.exception.OpenFailureException**

|                      |  |
|----------------------|--|
| Explanation          | This error occurs when the application client run time cannot read the Enterprise Archive (EAR) file.  |
| Possible causes      | The most likely cause of this error is that the system cannot find the EAR file cannot be found in the path specified on the <b>launchClient</b> command.  |
| Recommended response | Verify that the path and file name specified on the <b>launchClient</b> command are correct. If you are running on the Windows operating system and the path and file name are correct, use a short version of the path and file name (8 character file name and 3 character extension). |

**The launchClient command appears to hang and does not return to the command line when the client application has finished.**

|                      |  |
|----------------------|--|
| Explanation          | When running your application client using the <b>launchClient</b> command the WebSphere Application Server run time might need to display the security login dialog. To display this dialog the WebSphere Application Server run time creates an Abstract Window Toolkit (AWT) thread. When your application returns from its main method to the application client run time, the application client run time attempts to return to the operating system and end the Java Virtual Machine code. However, since there is an AWT thread, the Java Virtual Machine code will not end until <code>System.exit</code> is called. |
| Possible causes      | The Java Virtual Machine code does not end because there is an AWT thread. Java code requires that <code>System.exit()</code> be called to end AWT threads.  |
| Recommended response | <ul style="list-style-type: none"> <li>• Modify your application to call <code>System.exit(0)</code> as the last statement.</li> <li>• Use the <code>-CCexitVM=true</code> parameter when you call the <b>launchClient</b> command.</li> </ul>   |

---

## Chapter 8. Using Web services

### Before you begin

Decide whether you need a Web service solution for your application.

Use Web services when operating across a variety of platforms, including the Java 2, Enterprise Edition (J2EE) 1.3 and non-J2EE platforms. Web services benefit your e-business solution by integrating these enterprise systems, especially systems that have developed over a long period of time.

Using Web services makes most sense if your application's clients are non-J2EE applications, unless you have J2EE applications spread across the Web. It is recommended that you use J2EE technologies if all your clients are J2EE applications because performance can decrease when you use a Web service in a J2EE exclusive environment.

Because Web services are easily applied to existing applications and information technology assets, new solutions can be deployed quickly and recomposed to address new opportunities. As Web services become more popular, the pool of services grows, promoting development of more robust models of just-in-time application and business integration over the Internet.

To use Web services applications with WebSphere Application Server:

### Steps for this task

1. Plan to use Web services.

If you are not yet using Version 5.0.2 see Using Apache SOAP Web services in Version 5 and 5.0.1. Although the documentation remains in this release, support for Apache SOAP 2.x is deprecated and the function will be removed in a future release.

2. Migrate existing Web services.
3. Develop Web services.
4. Assemble Web services.
5. Deploy Web services.
6. Secure Web services.
7. Tune Web services.
8. Troubleshoot Web services.

### Usage scenario

The following is an example of how a business might use Web services.

The owner of a flower shop wants to start receiving orders from customers through the Web. She starts her venture by finding wholesale flower suppliers, pricing their product, and completing contracts for future flower orders.

Using Web services, the flower shop owner can find wholesale flower suppliers. One way she finds new suppliers is to use a Universal Description, Discovery, and

Integration (UDDI) registry to search for potential suppliers. She chooses the suppliers and the registry sends back information on how to contact the flower distributors that meet her criteria.

The flower shop owner can request price lists from each of the suppliers by obtaining a Web Services Description Language (WSDL) file for each potential supplier. The WSDL can be downloaded from the supplier's Web page, received through email, or retrieved from the supplier's UDDI registry entry.

The WSDL describes the procedure call. When using WebSphere Application Server, the procedure call is a Java API for XML-based remote procedure call (JAX-RPC), which helps her get price lists. The WSDL file also specifies the Universal Resource Locator (URL) where the request is to be sent.

The flower shop owner now has to compare the prices she received back from each supplier, decide which suppliers she is going to do business with, and make arrangements for future orders to be filled. The ground work has been laid for the flower shop to sell merchandise through the Web by using Web services to communicate with suppliers for the best prices and complete the ordering processes. The merchandise price lists need to be published to her Web site and she needs to provide a mechanism for customers to order flowers.

The flower supplier's Web services clients are deployed on the flower shop server. When a customer makes a transaction to purchase flowers through the Web, the order is sent to the supplier through JAX-RPC. The supplier responds by sending a confirmation with the order number and shipping date. The suppliers maintain the inventory and the flower shop owner handles billing and customer order management.

Similarly, the flower shop catalog can be composed automatically from the catalogs of all the suppliers. If the supplier ships directly to the customer, the order tracking inquiries can pass directly to the supplier's order tracking system. Web services can also be used by the supplier to send invoices for orders and by the flower shop to pay the supplier's invoices. Processes that previously required forms to be filled out manually, and faxed or mailed, can now be done automatically, saving labor costs for both the flower shop and the supplier.

Using Web services is beneficial because a much larger inventory is made available to the flower shop. There is no merchandise maintenance overhead, but the flower shop can offer their customers products that they otherwise might not have. Selling flowers through the Web increases capital for the flower shop without overhead of another store or money invested into additional product.

#### What to do next

After you decide to use Web services, design the model that best fits your e-business solution. You can learn more about Web services models in Planning to use Web services.

---

## Web services

*Web services* are self-contained, modular applications that you can describe, publish, locate, and invoke over a network.

A typical Web services scenario is a business application requesting a service from a given URL using Simple Object Access Protocol (SOAP) over a HyperText



Transport Protocol (HTTP) or Java Messaging Service (JMS) transport. The service receives the request, processes it, and returns a response. Examples of a simple Web service include weather reports and stock quotes. The method call is synchronous, that is, it waits until the result is available. Transaction Web services, supporting quotes, business-to-business (B2B) or business-to-client (B2C) operations include airline reservations or purchase orders.

Web services used by business operations are more complex, processing multiple requests or transactions at the same time. The time required to return the response depends on the complexity of the routing, therefore, the response is sent as an operation that is separate from the request.

Web services reflect a new, service-oriented approach to programming. This approach is based on the idea of building applications by discovering and implementing network-available services, or by invoking available applications to accomplish some task. This service-oriented approach is independent of specific programming languages or operating systems. Web services delivers interoperability; the ability for components created in different programming languages to work together as if they were created using the same language. Web services rely on existing transport technologies, such as HTTP, and standard data encoding techniques, such as Extensible Markup Language (XML), for invoking the implementation.

The key components of a Web service are:

- Simple Object Access Protocol (SOAP)
- Web Services Description Language (WSDL)
- Universal Description, Discover and Integration Protocol (UDDI)

There are four locations where a Java Web services client can reside:

- As an unmanaged stand-alone Java application
- Within an application client container
- As a Java bean or a servlet running in a Web container that is acting as a client
- As a Java bean or enterprise bean running in an enterprise Java bean (EJB) container that is acting as a client

You can review the Web services client programming model in the Web services for J2EE specification available through [Web services: Resources for learning](#). The programming model is similar to the EJB client programming model. There is a remote interface that the client uses to interact with the service. A Java Naming and Directory Interface (JNDI) lookup method can locate the service for a client running in a Web container, EJB container, or client container. The client obtains a stub that implements the remote interface and makes calls to invoke operations on the remote service.

---

## Planning to use Web services

### Before you begin

Plan your use of Web services and design an approach that best fits your e-business solution.

To plan to use Web services:

### Steps for this task

1. Design Web services to fit your e-business solution.  
Consider what you want to accomplish by using Web services, how Web services fit into your current topology, applications and programming model. Decide how the Web services process requests on the server and how the clients manage and use the Web service.  
Design your Web services for reliability, availability, manageability and security. For example, you want your Web services to process a transaction in a reasonable time at all hours of the day and provide users with good security characteristics, such as authentication for buyers. Planning to use Web services to work with WebSphere Application Server helps to meet these requirements.  
To support Web services, extend WebSphere Application Server to support Web services standards. For interoperable Web services running on platforms supplied by multiple vendors, standards are essential.
2. Install WebSphere Application Server.
3. Review Web services Samples.
4. Set up a Web services development environment.

#### What to do next

Develop a Web service.

## Setting up a Web services development environment

### Before you begin

WebSphere Application Server provides command-line tools to develop Web services clients and implementations. This topic describes how to set up a development environment to use these command-line tools.

WebSphere Application Server also includes the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>) that can be used for Web services assembly and deployment operations.

WebSphere Studio Application Developer Version 5.1 has GUI-based development tools to develop Web services that integrate with WebSphere Application Server 5.0.2.

Before you can set up a Web services development environment within WebSphere Application Server, you must install WebSphere Application Server.

To set up a Web services development environment:

### Steps for this task

1. Configure the path.

You can add the bin directory to your path by typing:

```
set PATH=install_root\bin;%PATH%
```

for Windows platforms, or

```
export PATH=install_root/bin:$PATH
```

for UNIX.

To use Java commands from the command shell, add the Java bin directory to your path by typing:

```
set PATH=install_root\java\bin;%PATH%
```

for Windows platforms, or

```
export PATH=install_root/java/bin:$PATH
```

for UNIX.Where:

*install\_root* is the fully qualified path to the WebSphere Application Server installation directory. The examples in this document assume the path includes both of these directories.

## 2. Configure the class path.

You must configure your CLASSPATH environment variable appropriately to compile your classes and execute Web services clients.

To configure the CLASSPATH, use the files listed in the following table. You can add the files from the WebSphere run time to your CLASSPATH, depending on the requirements of your code. Typically, only the *j2ee.jar* file is required to compile Web services implementations, while the other files are required to compile and run clients.

| File (Use forward slashes for UNIX)              | Contents   |
|--|--|
| <i>install_root</i> \lib\j2ee.jar                | J2EE classes, including Extendable Markup Language (XML) processing APIs |
| <i>install_root</i> \lib\xerces.jar              | XML manipulation classes needed by the client run time                   |
| <i>install_root</i> \lib\commons-logging-api.jar | Client run time logging support  |
| <i>install_root</i> \lib\wsdl4j.jar              | Client run time Web Services Description Language (WSDL) file utilities  |
| <i>install_root</i> \lib\webservices.jar         | Client run time  |
| <i>install_root</i> \lib\qname.jar               | XML classes  |
| <i>install_root</i> \lib\commons-discovery.jar   | Client run time support  |

### What to do next

Develop Web services.

---

## Migrating Apache SOAP Web services to Web services for J2EE

### Before you begin

If you have used the Apache SOAP support to create Web services client applications in WebSphere Application Server Versions 4.0 or 5.0, and want use Web services for J2EE, also known as JSR 109, you need to migrate your Version 4.0 and 5.0 client applications.

To migrate these client applications to the JSR 109 Web services standards:

### Steps for this task

#### 1. Plan your migration strategy.

There are two ways you can port an Apache SOAP client to Java API for XML-based RPC (JAX-RPC) Web services client:

- If you have, or can create, a Web Services Description Language (WSDL) document for the service, consider using the **WSDL2Java** command tool to

generate bindings for the Web service. It is more work to adapt an Apache SOAP client to use the generated JAX-RPC bindings, but the resulting client code is more robust and easier to maintain. To follow this path, see *Developing a Web services client*.

- If you do not have a WSDL document for the service, do not expect the service to change, and you want to port the Apache SOAP client with a minimal work, you can convert the code to use the JAX-RPC dynamic invocation interface (DII), which is similar to the Apache SOAP APIs. The DII APIs do not use WSDL or generated bindings.

You should be aware that since JAX-RPC does not specify a framework for user-written serializers, the JAX-RPC does not support the use of custom serializers. If your application cannot conform to the default mapping between Java, WSDL, and XML supported by WebSphere Application Server, you should not attempt to migrate the application.

The remainder of this topic assumes that you have decided to use the JAX-RPC DII APIs.

## 2. Review the GetQuote sample.

There is a Web services migration sample in the Samples Gallery. This sample is located in the `GetQuote.java` file, originally written for Apache SOAP, and includes an explanation about the changes needed to migrate to the JAX-RPC DII interfaces.

## 3. Convert the client application from Apache SOAP to JAX-RPC DII

The Apache SOAP API and JAX-RPC DII API structures are similar. You can instantiate and configure a call object, set up the parameters, invoke the operation, and process the result in both.

You can create a generic instance of a Service object with

```
javax.xml.rpc.Service service =  
    ServiceFactory.newInstance().createService(new QName(""));
```

in JAX-RPC.

### a. Create the call object.

An instance of the call object is created by

```
org.apache.soap.rpc.Call call = new org.apache.soap.rpc.Call ()
```

in Apache SOAP.

An instance of the call object is created by

```
java.xml.rpc.Call call = service.createCall();
```

in JAX-RPC.

### b. Set the endpoint URI.

The target URI for the operation is passed as a parameter to

```
call.invoke: call.invoke("http://...", "");
```

in Apache SOAP.

The `setTargetEndpointAddress` method is used as a parameter to

```
parametcall.setTargetEndpointAddress("http://...");
```

in JAX-RPC.

Apache SOAP has a `setTargetObjectURI` method on the call object that contains routing information for the request. JAX-RPC has no equivalent method. The information in the `targetObjectURI` is included in the `targetEndpointURI` for JAX-RPC.

- c. Set the operation name.

The operation name is configured on the call object by

```
call.setMethodName("opName");
```

in Apache SOAP.

The `setOperationName` method, which accepts a `QName` instead of a `String` parameter, is used in JAX-RPC as follows:

```
call.setOperationName(new javax.xml.namespace.QName("namespace", "opName"));
```

- d. Set the encoding style.

The encoding style is configured on the call object by

```
call.setEncodingStyleURI(org.apache.soap.Constants.NS_URI_SOAP_ENC);
```

in Apache SOAP.

The encoding style is set by a property of the call object

```
call.setProperty(javax.xml.rpc.Call.ENCODINGSTYLE_URI_PROPERTY,
    "http://schemas.xmlsoap.org/soap/encoding/");
```

in JAX-RPC.

- e. Declare the parameters and set the parameter values.

Apache SOAP parameter types and values are described by parameter instances, which are collected into a `Vector` and set on the call object before the call, for example:

```
Vector params = new Vector ();
params.addElement(
    new org.apache.soap.rpc.Parameter(name, type, value, encodingURI));
// repeat for additional parameters... call.setParams (params);
```

For JAX-RPC, the call object is configured with parameter names and types without providing their values, for example:

```
call.addParameter(name, xmlType, mode);
// repeat for additional parameters call.setReturnType(type);
```

Where

- *name* (type `java.lang.String`) is the name of the parameter
- *xmlType* (type `javax.xml.namespace.QName`) is the XML type of the parameter
- *mode* (type `javax.xml.rpc.ParameterMode`) the mode of the parameter, for example, IN, OUT, or INOUT

- f. Make the call.

The operation is invoked on the call object by

```
org.apache.soap.Response resp = call.invoke(endpointURI, "");
```

in Apache SOAP.

The parameter values are collected into an array and passed to `call.invoke` as follows:

```
Object resp = call.invoke(new Object[] {parm1, parm2,...});
```

in JAX-RPC.

g. Check for faults.

You can check for a SOAP fault on the invocation by checking the Response:

```
if resp.generatedFault then {  
  org.apache.soap.Fault f = resp.getFault();  
  f.getFaultCode();  
  f.getFaultString(); }  
}
```

in Apache SOAP.

A `java.rmi.RemoteException` is thrown in JAX-RPC if a SOAP fault occurs on the invocation.

```
try  
  ... call.invoke(...)  
catch (java.rmi.RemoteException) ...
```

h. Retrieve the result.

In Apache SOAP, if the invocation was successful and returns a result, it can be retrieved from the Response object:

```
Parameter result = resp.getReturnValue(); return result.getValue();
```

In JAX-RPC, the result of invoke is the returned object when no exception is thrown:

```
Object result = call.invoke(...);  
...  
return result;
```

### What to do next

Develop a Web services client.

Test the Web services-enabled clients.

---

## Developing Web services

### Before you begin

WebSphere Application Server uses Web services standards developed for the Java language under the Java Community Process (JCP). These standards include the Java API for XML-based remote procedure call (JAX-RPC (JSR-101)) and Web services for J2EE (JSR-109).

The JAX-RPC standard covers the programming model and bindings for using Web Services Description Language (WSDL) for Web services in the Java language. The Web services standard for J2EE covers the use of JAX-RPC in a J2EE environment, as well as the deployment of Web services implementations in a J2EE server. Both standards are part of the J2EE 1.4 release.

For more information on JAX-RPC, JSR-109, tutorials and other Web services and J2EE information, see [Web services: Resources for learning](#).

You can also use the WebSphere Studio Application Developer Version 5.1 graphical user interface development tools to develop Web services that integrate with WebSphere Application Server.

Before you develop Web services, see the topic [Setting up a Web services development environment](#).

You can develop a J2EE Web service application in one of four ways:

#### Steps for this task

1. Develop a Web service using a Java bean.
2. Develop a Web service using a stateless session enterprise bean.
3. Develop a Web service with an existing WSDL file using a Java bean.
4. Develop a Web service with an existing WSDL file using a stateless session enterprise bean.

#### What to do next

Assemble the Web service.

## **Developing a Web service using a Java bean**

### Before you begin

Set up a Web service development environment.

To develop a Web service using a Java bean:

#### Steps for this task

1. Access an existing Java bean Web archive (WAR) file.
2. Develop a Service Endpoint Interface.
3. Develop a Web Services Description Language (WSDL) file.
4. Develop Web service deployment descriptor templates when starting with WSDL file.
  - a. Configure the webservices.xml deployment descriptor with command-line tools.
  - b. Configure the webservices.xml deployment descriptor with the Assembly Toolkit.
  - c. Configure the ibm-webservices-bnd.xmi deployment descriptor with command-line tools.
  - d. Configure the ibm-webservices-bnd.xmi deployment descriptor with the Assembly Toolkit.
5. Assemble a Web services-enabled WAR file when starting from Java.
6. Assemble a Web services-enabled enterprise archive (EAR) file.
7. Deploy the EAR file into WebSphere Application Server.

#### What to do next

Assemble the Web service.

## **Developing a Service Endpoint Interface**

### Before you begin

The Service Endpoint Interface defines the Web services methods. The Web service implementation, whether an enterprise bean or a Java bean, must implement methods having the same signature as the methods on the Service Endpoint Interface. There are a number of restrictions on which types to use as parameters

and results of Service Endpoint Interface methods. These restrictions are documented in the Java API for XML remote procedure call (JAX-RPC) specification, which is available through Web services: Resources for learning.

It is easiest to develop the Service Endpoint Interface from an enterprise JavaBean (EJB) remote interface when the Web service implementation is an enterprise bean. When the Web service implementation is a Java bean, the bean or an interface it implements can be a basis for the Service Endpoint Interface.

To develop a Service Endpoint Interface:

#### Steps for this task

1. Create a Java interface containing the methods to include in the Service Endpoint Interface.

The interface should extend the `java.rmi.Remote` interface. Each method throws the exception, `java.rmi.RemoteException`. If you start with an existing Java interface, remove any methods that do not conform to JAX-RPC.

2. Compile the interface.

You need, `install_root\lib\j2ee.jar`, in your CLASSPATH to compile the interface.

#### Results

A Service Endpoint Interface which you can use to develop a Web service.

#### Usage scenario

This example uses a Java interface called `AddressBook`. The following example depicts the `AddressBook` interface:

```
package addr;
public interface AddressBook extends java.rmi.Remote {
    /**
     * Retrieve an entry from the AddressBook.
     *
     * @param name the name of the entry to look up.
     * @return the AddressBook entry matching name or null if none.
     * @throws java.rmi.RemoteException if communications failure.
     */
    public addr.Address getAddressFromName(java.lang.String name)
        throws java.rmi.RemoteException;
}
```

You use the `AddressBook` Java interface to create the Service Endpoint Interface:

1. Begin with the remote interface, `AddressBook_RI.java`:

```
package addr;
import java.rmi.RemoteException;
import javax.ejb.EJBObject;
public interface AddressBook_RI extends EJBObject {
    /**
     * Retrieve an entry from the AddressBook.
     *
     * @param name the name of the entry to look up.
     * @return the AddressBook entry matching name or null if none.
     * @throws java.rmi.RemoteException if communications failure.
     */
    public addr.Address getAddressFromName(java.lang.String name)
        throws java.rmi.RemoteException;
}
```



2. Make a copy of the remote interface named `AddressBook.java` and use it as a template for the Service Endpoint Interface.
3. Change the interface to extend the `java.rmi.Remote` interface, instead of the `EJBObject` interface.
4. Compile the interface.

#### What to do next

Use the Service Endpoint Interface to Develop a Web Services Description Language (WSDL) file.

### **Developing a Web Services Description Language file**

#### Before you begin

Develop a Service Endpoint Interface. You must configure the Service Endpoint Interface class file and the classes in your CLASSPATH.

You need a Web Services Description Language (WSDL) file to use Web services. You can develop your own WSDL file or get one from a Web service provider through E-mail, downloading or through a Uniform Resource Locator (URL). This documentation assumes you are creating your own.

To develop a WSDL file:

#### Steps for this task

1. Run the **Java2WSDL** *seiInterface* command.
  - a. Move the WSDL file to the `META-INF/wsdl` subdirectory if you are using enterprise JavaBeans (EJB).
  - b. Move the WSDL file to the `WEB-INF/wsdl` subdirectory if you are using a Java bean.

A WSDL file named *seiInterface.wsdl* is created.

2. Edit the generated WSDL file and inspect the part names.

The WSDL parts have names like `arg_0_0`. Modify the WSDL file to use the actual names of the Java parameters.
3. (Optional) Use the **Java2WSDL** command tool to generate the correct part names of WSDL file.

You can automatically generate and set the correct part names by using the **Java2WSDL** command tool. Generating and setting the part names is done by providing additional information to the **Java2WSDL** command in the form of a Java implementation class that implements the same methods as the Service Endpoint Interface and is compiled with debug information on (**javac -g**). Parameter names are stored in the `.class` file with the debug information. If your implementation class was compiled with debug on, you can use the **Java2WSDL -implClass seiImpl seiInterface** command to generate a WSDL file having the proper part names.

#### Results

A WSDL file that defines the Web service described by the Service Endpoint Interface.

#### Usage scenario

This example uses a JAR file named `AddressBook.jar` containing a class named `AddressBook.class` file.

You must add the `AddressBook.jar` file to your CLASSPATH to create the WSDL file. The JAR file contains an enterprise JavaBean (EJB) implementation class that was compiled with debugging information on. Run the `Java2WSDL -implClass addr.AddressBookBean addr.AddressBook` command to create a WSDL file named `AddressBook.wsdl`.

#### What to do next

Develop Web services deployment descriptor templates from a WSDL file.

### **Publishing Web Services Description Language files**

#### Before you begin

You need an enterprise application, also known as an enterprise archive (EAR) file, that contains a Web services-enabled module and has been configured and enabled for Web services. See [Deploying Web services](#).

The Web Services Description Language (WSDL) files for each Web services-enabled module are published to the file system location you specify. You can provide these WSDL files to clients that want to invoke your Web services.

You can publish WSDL files for the deployed EAR file in one of three ways:

#### Steps for this task

1. Publish a WSDL file with the administrative console.
2. Publish a WSDL file with the `wsadmin` command tool.
3. Publish a WSDL file through a URL.

### **Publishing Web Services Description Language files with the administrative console**

#### Before you begin

When publishing Web Services Description Language (WSDL) files with the administrative console, you can specify default or custom HTTP URL prefixes. You can also specify a Java Message Service (JMS) URL prefix.

To publish a WSDL file with the administrative console:

#### Steps for this task

1. Open the administrative console.
2. Click **Applications > Enterprise Applications > *application***. Under Additional Properties, click **Publish WSDL** which brings you to the **Publish WSDL files for Web Services** panel.
3. Specify the default URL prefixes for the Web service.
  - a. Select **HTTP URL prefix**.
  - b. Select an entry from the drop down list.

If you have multiple application modules, select the application module's checkbox on the module table.
  - c. Click **Apply**.

The URL prefix is copied to the selected module HTTP URL prefix field.

- d. Click **OK**.
  - e. Click the exported *WSDL\_zip\_file* listed on the **Export WSDL Zip file** panel.
  - f. Download the zip file.  
Follow your browser's instructions to download the zip file.
4. Specify custom URL prefixes for the Web service.
    - a. Select **Custom HTTP URL prefix**.
    - b. Type the name of the URL prefix in the **Custom HTTP URL prefix** field.  
The entry must be of the form `http|https://<host_name>:<port_number>`.  
For example:  
`http://myHost:999`  
  
If you have multiple application modules, select the application module's checkbox on the module table.
    - c. Click **Apply**.  
The URL prefix is copied to the selected module HTTP URL prefix field.
    - d. Click **OK**.
    - e. Click the exported *WSDL\_zip\_file* listed on the **Export WSDL Zip file** panel.
    - f. Download the zip file.  
Follow your browser's instructions to download the zip file.
  5. Specify a JMS URL prefix.
    - a. Select the application module.
    - b. Type the JMS URL prefix into the **JMS URL prefix** field.  
The entry must be of the form: `jms:[queue|topic]?destination=<queue or topic_jndi_name>&connectionFactory=<connection_factory_jndi_name>`. For example:  
`jms:/queue?destination=jms/Q1&connectionFactory=jms/QCF1`
    - c. Click **OK**.
    - d. Click the exported *WSDL\_zip\_file* listed on the **Export WSDL Zip file** panel.
    - e. Download the zip file.  
Follow your browser's instructions to download the zip file.

## **Publishing Web Services Description Language files with the *wsadmin* command tool**

### Before you begin

The Web Services Description Language (WSDL) files in each Web services-enabled module are published to the file system location you specify. You can provide these WSDL files to the clients that want to invoke your Web services.

The scripting client (**wsadmin**) can publish the WSDL files in either local, for example, `-conntype NONE`, or remote mode. However, in local mode, the target application should be located at the same node where the **wsadmin** command is invoked.

The steps below assume that the application server is running.

To publish a WSDL file with the **wsadmin** command:

### Steps for this task

1. From a command prompt, start *install\_root\bin\wsadmin* if you are using Windows or *install\_root/bin/wsadmin* if you are using UNIX.
2. At the **wsadmin** command prompt, enter one of the two commands:
  - **\$AdminApp publishWSDL *app\_Name path\_Name***
  - **\$AdminApp publishWSDL *app\_Name path\_Name* {{module {{binding *url-prefix*}}}}**

Where

- *app\_Name* is the application name
- *path\_Name* is the absolute path to the zip file that will contain the published WSDL files. The zip file is saved in the server side, therefore, if the server is running on a different machine, you need to obtain the zip file from that machine. The directory structure of the resulting zip file is based on the following information:

```

Application file name
  module file name
    META-INF/ or WEB-INF/
      wsdl/
        WSDL file name
  
```

See the usage scenario for an example of this directory structure.

- *binding* is either http or jms (both are in lower case)
- *url-prefix* is the partial SOAP address for the associated SOAP binding. For an HTTP binding the form is `http://host:port/` or `https://host:port.` For Java Message Service (JMS) bindings, the form is `jms:/queue?destination=dest&connectionFactory=cf` or `jms:/topic?destination=dest&connectionFactory=cf` The **\$AdminApp publishWSDL *app\_Name path\_Name*** command updates the WSDL Simple Object Access Protocol (SOAP) address prefixes using the default values. If you do not want to update the WSDL SOAP address prefixes, use the other command, instead of the default values. The **\$AdminApp publishWSDL *app\_Name path\_Name* {{module {{binding *url-prefix*}}}}** command allows you to customize the WSDL SOAP address for each module. You can specify a different address prefix for each SOAP binding. Results The WSDL files from Web services are published to a specified zip file. You can hand the zip file to the client and the client can use the published WSDL files to create a Web services client that accesses the deployed service. Usage scenario The command to publish WSDL files for a Web service named `WebServicesSamples` could be **\$AdminApp publishWSDL WebServicesSamples *c:/temp/samplesWsd1.zip*** or **\$AdminApp publishWSDL WebServicesSamples *c:/temp/sampleswsdl.zip* { {AddressBookJ2WB.war { {http *http://localhost:9080*}} } {StockQuote.jar { {http *https://localhost:9443*}} } }** The directory structure for this created zip files is `WebServicesSamples.ear/StockQuote.jar/META-INF/wsdl/StockQuoteFetcher.wsdl`  
`WebServicesSamples.ear/AddressBookW2JE.jar/META-INF/wsdl/AddressBookW2JE.wsdl`  
`WebServicesSamples.ear/AddressBookJ2WE.jar/META-INF/wsdl/AddressBookJ2WE.wsdl`  
`WebServicesSamples.ear/AddressBookJ2WB.war/WEB-INF/wsdl/AddressBookJ2WB.wsdl`  
`WebServicesSamples.ear/AddressBookW2JB.war/WEB-INF/wsdl/AddressBookW2JB.wsdl`

## Publishing Web Services Description Language files through a URL

### Before you begin

The Web services-enabled application should be installed and running.

The files referenced by the `<wsdl-file>` element in the `webservices.xml` file can or cannot import other Web Services Description Language (WSDL) or XSD files. Typically, all WSDL or XSD files are originally placed into the `META-INF/wsdl` directory when using enterprise JavaBeans (EJBs) or the `WEB-INF/wsdl` directory when using Java beans. If your WSDL or XSD files are not placed in one of these directories, the file referenced by the `<wsdl-file>` and its imported files are located at the same directory and copied to the `wsdl/` directory for publishing purposes.

**Note:** EJB-based Web service applications must have an HTTP router or a Web module. The publishing URL supported is of HTTP type.

To publish a WSDL file through a URL:

### Steps for this task

1. Retrieve the outer-most WSDL file.

The outer-most WSDL file is the WSDL file defined by the `<wsdl-file>` element in the `webservices.xml` file.

Each Web service has an endpoint address, like

`http://example.com/services/stockquote`. You can retrieve the outer-most WSDL file (defined by the `<wsdl-file>` element within the `webservices.xml` file) by appending the string `"/wsdl"` or `"/wsdl/"` to the endpoint address, for example, `http://example.com/services/stockquote/wsdl`.

2. Retrieve the imported WSDL files.

When the outer-most WSDL file imports other WSDL or XSD files, these imported files can be retrieved by appending the relative path to the URL, which is used to retrieve the outer-most WSDL file. This is also true for WSDL files that import other files. This process is similar to typical HTTP protocol. If an HTML document contains a hyperlink to other documents, the relative path is appended to create the URL to access the hyperlinked documents.

### Usage scenario

Suppose you have an application with the following directory structure:

```
<module-root>/
META-INF/
WEB-INF/
webservices.xml
/* define Foo service, the <wsdl-file> element points to "/wsdl/fooImpl.wsdl" */
web.xml
ibm-webservices-bnd.xml
<jaxrpc-mapping-file>
wsdl/
fooImpl.wsdl/* importing foo.wsdl which is an interface wsdl */
foo.wsdl /* importing type definition for the interface */
fooTypes.xsd
```

If the SOAP address for the foo service is

`http://examples.com:9080/services/foo`, the simple way to retrieve the foo's outer-most WSDL, is with the following form

`http://examples.com:9090/services/foo/wsdl` or

`http://examples.com:9090/services/foo/wsd1/`. The URL is redirected to `http://examples.com:9090/services/foo/wsd1/fooImpl.wsd1`, where `fooImpl.wsd1` is the name of the outer-most WSDL file.

Since the `fooImpl.wsd1` file has the import `<import namespace="http://examples.com/foo" location="a/b/foo.wsd1">`, use the URL `http://examples.com:9090/services/foo/wsd1/a/b/foo.wsd1` to obtain the `foo.wsd1` file.

## **Multipart Web Services Description Language file best practices**

WebSphere Application Server supports deployment of Web services using a multipart Web Services Description Language (WSDL) file. That is, WSDL files import other WSDL files when the WSDL file listed in the `<wsdl-file>` element of the `webservices.xml` deployment descriptor contains all `<wsdl:service>` and `<wsdl:port>` elements. The WSDL file is divided into an implementation WSDL and an interface WSDL.

The `<wsdl:import>` element indicates a reference to another WSDL file. If the `<wsdl:import>` element location attribute does not contain a URL, that is, it contains only a file name, and does not begin with `http://`, `https://` or `file://`, the imported file must be located in the same directory and must not contain a relative path component. For example, if `META-INF/A_Impl.wsd1` is in your module and contains the import statement `<wsdl:import="A.wsd1" namespace="...">`, the file, `A.wsd1` must also be located in the module `META-INF` directory.

It is recommended that all WSDL files be placed in either the `META-INF/wsd1` directory, if you are using enterprise JavaBeans (EJBs), or the `WEB-INF/wsd1` directory, if you are using Java beans, even if there are relative imports within the WSDL files. Otherwise, there are implications when the WSDL publication is involved with `<location="./interfaces/A_Interface.wsd1" namespace="...">`. Using a path like this fails due to the presence of the relative path, regardless of whether the file is located at that path or not. If the location is a URL, it must be readable at both deployment and server startup.

### **WSDL publication**

The files located in the `META-INF/wsd1` or `WEB-INF/wsd1` directory can be published through either a URL or file, including WSDL or XSD files. For example, if the file referenced in the `<wsdl:file>` element of the `webservices.xml` deployment descriptor is located in the `META-INF/wsd1` or `WEB-INF/wsd1` directory, it is publishable. If the files imported by the `<wsdl:file>` are located in the `wsdl/` directory or its subdirectory, they are publishable.

If the WSDL file referenced by the `<wsdl:file>` element is located in a directory other than `wsdl/`, or its subdirectories, the file and its imported files, either WSDL or XSD files, which are in the same directory, are copied to the `wsdl/` directory without modification when the application is installed. These types of files can also be published.

If the `<wsdl:file>` imports a file located in a different directory, the file is not copied to the `wsdl/` directory and not available for publishing.

## **Developing Web services deployment descriptor templates from a Web Services Description Language file**

Before you begin

To develop the deployment descriptor templates from a Web Services Description Language (WSDL) file, you must obtain the Uniform Resource Locator (URL) of the WSDL file to use.

If it is a local file and you are running the Windows platform, the URL looks like this: *file:drive:\path\file\_name.wsdl*. If you are using the UNIX platform, the URL looks like this: *file:/path/file\_name.wsdl*. You can also specify local files using the absolute or relative file system path.

The client deployment descriptors, *webservicesclient.xml* and *ibm-webservicesclient-bnd.xmi*, are generated when the **WSDL2Java -role** argument specifies a client role. The server deployment descriptors, *webservices.xml* and *ibm-webservices-bnd.xmi* are generated when the **-role** argument is a server role.

To develop deployment descriptor templates:

#### Steps for this task

1. Run the **WSDL2Java -verbose -role develop-client -container *type* -genJava No *wsdlURL*** command or **WSDL2Java -verbose -role develop-server -container *type* -genJava No *wsdlURL*** command to generate the client and server deployment descriptor templates into the META-INF or WEB-INF subdirectory.

Where:

*type* is EJB for an enterprise JavaBean (EJB)-based implementation or Web for a Java bean-based implementation. *wsdlURL* is the URL of the WSDL file to use.

The **-container** option that you choose to run determines to which subdirectory the templates are generated.

The Java API for XML-based remote procedure call (JAX-RPC) mapping file is needed for both server and client use, and is always generated. When the **WSDL2Java -container EJB** argument is specified, all deployment descriptors and the JAX-RPC mapping file are generated into the META-INF subdirectory of the output directory. When **-container Web** argument is specified, all deployment descriptors and the JAX-RPC mapping file are generated into the WEB-INF subdirectory of the output directory. To generate the deployment descriptors only, and not any Java classes, specify the **-genJava No** argument with the **WSDL2Java** command tool.

If the **-verbose** option is specified, a list of all generated files displays when the command runs.

#### Results

Deployment descriptor templates that are required to implement or use a Web service.

#### Usage scenario

The following example uses a WSDL file named *AddressBookJ2WB.wsdl*:

1. Generate the template files:
  - **WSDL2Java -verbose -role develop-client -container Web -genJava No META-INF\AddressBookJ2WB.wsdl**
  - **WSDL2Java -verbose -role develop-server -container Ejb -genJava No META-INF\AddressBookJ2WB.wsdl**

The deployment descriptor templates are generated into the WEB-INF for client and META-INF for EJB server subdirectories as follows:

```
Parsing XML file: META-INF/AddressBookJ2WB.wsdl
Generating: WEB-INF\webservicesclient.xml
Generating: WEB-INF\ibm-webservicesclient-bnd.xmi
Generating: WEB-INF\AddressBookJ2WB_mapping.xml
Generating: META-INF\webservices.xml
Generating: META-INF\ibm-webservices-bnd.xmi
Generating: META-INF\AddressBookJ2WB_mapping.xml
```

### What to do next

Configure the deployment descriptors.

## **Configuring the webservices.xml deployment descriptor with command-line tools**

### Before you begin

Develop implementation templates and bindings for the webservices.xml deployment descriptor from the Web Services Description Language (WSDL) file.

This topic explains how to configure the webservices.xml deployment descriptor with command-line tools. You can also configure the webservices.xml deployment descriptor using the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>). See *Configuring the webservices.xml deployment descriptor with the Assembly Toolkit*.

To configure the webservices.xml deployment descriptor using command-line tools:

### Steps for this task

1. Open the webservices.xml deployment descriptor in a text editor.
2. Inspect and complete the file elements.

The file elements contain default values. The following steps identify each element and its contents. Inspect the values for accuracy and fill in default values that do not exist. Elements that do not have default values are flagged with "???" in the XML file. To set the elements:

- a. Locate the <webservice-description-name> attribute.

This name must be unique within the XML file if there are multiple webservice-description elements. This name is used as part of the directory structure for naming published WSDL files when publishing files during deployment.

- b. Locate the <wsdl-file> element.

Change the tag element from the full path within the Java archive (JAR) or Web archive (WAR) file to the WSDL file defining the Web service being implemented. By convention, the WSDL file is placed in the META-INF/wsdl subdirectory for EJB modules and the WEB-INF/wsdl subdirectory for Web modules.

- c. Locate the <jaxrpc-mapping-file> element.

Set this element to the full path within the JAR or WAR file to the generated Java API for XML-based remote procedure call (JAX-RPC) mapping file. By convention, the mapping file is placed in the META-INF subdirectory for EJB modules and the WEB-INF subdirectory for Web modules.

- d. Locate the <port-component-name> element.



Change the element to a string that uniquely identifies the port within the module.

- e. Locate the <wsdl-port> element.

This identifies the port in the WSDL file that corresponds to this deployment descriptor port. The <namespaceURI> subelement indicates the namespace portion of the WSDL port name. Set this element to the namespace of the matching port in the WSDL file identified by the <wsdl-file> element. This namespace is typically the target namespace in the WSDL file. The <localpart> subelement indicates the local portion of the WSDL port name. Set this element to the name of the matching WSDL port.

- f. Locate the <service-endpoint-interface> element.

Set this element to the class name, including package, of the Java interface that is the Service Endpoint Interface for the Web service.

- g. Locate the <ejb-link> element if the Web service implementation is an enterprise bean.

Set this element to the <ejb-name> element of ejb-jar.xml file for the corresponding enterprise bean. This link informs the Web service port component which enterprise bean implements it.

- h. Locate the <service-impl-bean> tag if the Web service implementation is a Java bean in a Web module.

Set the <servlet-link> element to reference the bean that implements the methods of the Service Endpoint Interface.

## Results

The webservices.xml deployment descriptor is configured for the Web service implementation module. To learn more about the assembly properties for the webservices.xml deployment descriptor see the Web services for J2EE specification available through Web services: Resources for learning.

## What to do next

Configure the ibm-webservices-bnd.xmi deployment descriptor.

## **Configuring the webservices.xml deployment descriptor with the Assembly Toolkit**

### Before you begin

Create an enterprise JavaBean (EJB) Java archive (JAR) file or Web archive (WAR) file containing webservices.xml.

This topic explains how to configure the webservices.xml deployment descriptor with the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>). For more information about completing tasks with the Assembly Toolkit, click **Help > Help** in the Assembly Toolkit graphical user interface (GUI).

You can also configure the webservices.xml deployment descriptor with command-line tools. See Configuring the webservices.xml deployment descriptor with command-line tools.

To configure the webservices.xml deployment descriptor with the Assembly Toolkit:

### Steps for this task

1. Start the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>).
2. Click **File > Import** to import the EJB JAR file or WAR file into the Assembly Toolkit.
3. Right-click the `webservices.xml` file in the navigator pane.  
You might have to switch to the Java perspective for the `webservices.xml` file to be visible in the navigator pane.
4. Select **Open**.  
The **Web Services** editor opens.
5. Expand the **Web service descriptions** section.
  - a. Select the service you want to configure.
6. Expand the **Web service description implementation details** section.
  - a. Verify the **Web service description name** field is unique among all the Web service descriptions in the editor.
  - b. Verify that the **WSDL file** field indicates there is an existing WSDL file in the module.  
This file, by convention, should be located in the `META-INF/wsdl` directory for an enterprise bean JAR file and in the `WEB-INF/wsdl` directory for a WAR file.
  - c. Verify the **JAX-RPC mapping file** field indicates an existing mapping file within the module.  
This file, by convention, should be located in the `META-INF` directory for an enterprise bean JAR file and in the `WEB-INF` directory for a WAR file.
7. Expand the **Port components** section.
  - a. Verify there are port component entries corresponding to the used WSDL ports in the **Port components** section.
8. Select a *port\_component* to open the editor for that port component.  
The **Port Components** editor opens.
9. Expand the **Port component implementation details** section.
  - a. Verify the **WSDL Port Namespace URL** and **WSDL Port Local part** fields are set to the namespace and local name of the corresponding port in the WSDL file.  
These fields are configured by the **WSDL2Java** command tool when the `webservices.xml` file is generated.
10. Verify the **Service endpoint interface** field names the fully qualified Service Endpoint Interface class.  
This field is configured by the **WSDL2Java** command when the `webservices.xml` file is generated.
11. Locate the **Service implementation bean** field.
  - a. Configure this field to indicate the EJB or servlet that implements the Web service. Configure by selecting **EJB link** for an enterprise bean module or **Servlet link** for a Web module. Use the drop down list in the **Service implementation bean** field to select the enterprise bean or servlet used to implement the Web service.  
The choices in the drop down menu come from the enterprise beans defined in the `ejb-jar.xml` file for an enterprise bean module or the servlets defined in the `web.xml` file for a Web module.

## Configuring the ibm-webservices-bnd.xmi deployment descriptor with command-line tools

### Before you begin

Develop implementation templates and bindings for the `ibm-webservices-bnd.xmi` from the Web Services Description Language (WSDL) file.

This topic explains how to configure bindings using command-line tools. You can also configure the bindings with the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>). See [Configuring the ibm-webservices-bnd.xmi deployment descriptor with the Assembly Toolkit](#)

To configure the `ibm-webservices-bnd.xmi` deployment descriptor with command-line tools:

### Steps for this task

1. Open the `ibm-webservices-bnd.xmi` deployment descriptor in a text editor.
2. Inspect and complete the file elements.

The file elements contain default values. The following steps identify each element and its contents. Inspect the values for accuracy.

To set the elements:

- a. Locate the `<wsdescBindings>` element.

Set the `wsDescNameLink` attribute to the value of the `<webservice-description-name>` element in the `webservices.xml` deployment descriptor.

Set the `pcNameLink` attribute to the value of the `<port-component-name>` of the corresponding port in the `webservices.xml` deployment descriptor.

### Results

The `ibm-webservices-bnd.xmi` deployment descriptor is configured for the Web service implementation module.

### What to do next

Do one of the following based on whether your implementation is an EJB Java archive (JAR) file or Web module Web archive (WAR) file:

- Assemble a Web services-enabled JAR file when starting from Java.
- Assemble a Web services-enabled JAR file when starting from WSDL.

## Configuring the ibm-webservices-bnd.xmi deployment descriptor with the Assembly Toolkit

### Before you begin

Develop implementation templates and bindings for the `ibm-webservices-bnd.xmi` from the Web Services Description Language (WSDL) file.

Do one of the following based on whether your implementation is an EJB Java archive (JAR) file or Web module Web archive (WAR) file:

- Assemble a Web services-enabled JAR file when starting from Java.
- Assemble a Web services-enabled JAR file when starting from WSDL.

This topic explains how to configure bindings using the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>). For more

information about completing tasks with the Assembly Toolkit, click **Help > Help** in the Assembly Toolkit graphical user interface.

You can also configure the bindings with command-line tools. See *Configuring the ibm-webservices-bnd.xmi deployment descriptor with command-line tools*.

To configure the `ibm-webservices-bnd.xmi` deployment descriptor with the Assembly Toolkit:

#### Steps for this task

1. Start the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>).
2. Click **File > Import** to import the EJB JAR file or WAR file into the Assembly Toolkit.
3. Right-click the `webservices.xml` file in the Navigator pane.  
You might have to switch to the Java perspective for the `webservices.xml` file to be visible in the Navigator pane.
4. Select **Open**.  
The **Web Services** editor opens.
5. Click the **Bindings** tab located at the bottom of the editor pane.  
The **Web Services Bindings** editor opens.
6. Verify the `wsdescNameLink` element settings.
  - a. Expand the **Web services description bindings** section. Verify that the value of the `<webservice-description-name>` element in the `webservices.xml` deployment descriptor is listed in the section.  
If the value is not listed:
    - b. (Optional) Click **Add**, choose the correct Web services name and click **OK**.  
You do not need to complete this step if you have verified that the correct Web services name is listed in the **Web Services Description Bindings** tab.
7. Verify the `pcnameLink` attribute settings.
  - a. Expand the **Web Service Description Bindings** section. Verify that the correct service is selected in the section.  
If the correct service is not listed:
    - b. (Optional) Expand **Port Component Binding**.  
Verify the correct Web services name is selected in the **Web Service Description Bindings** section.  
This selection is a prerequisite to creating a `pcnameLink` attribute.
    - c. In the **Port Component Binding** section, click **Add**.  
You need to make a selection in the **Web Service Description Bindings** section before you can create the port component binding in the **Port Component Binding** section. The **Port Component Bindings Dialog** opens.
    - d. Select the desired port from the drop down list in the **PC Name Link** field.
    - e. Click **OK**.
    - f. Click the **Binding Configurations** tab to view the bindings for your port.
    - g. **Save** the bindings file.
8. Click **File > Export** to export the JAR file, or continue using the Assembly Toolkit for configuration and assembly tasks.

#### Results

The `ibm-webservices-bnd.xml` deployment descriptor is configured for the Web service implementation module.

## Configuring the `webservicesclient.xml` deployment descriptor with command-line tools

Before you begin

You should have generated the template for `webservicesclient.xml` from the Web Services Description Language (WSDL) file.

If the default values set in the template are acceptable, you do not need to modify the file further.

If the Web services client is in an enterprise JavaBean (EJB) Java archive (JAR) file, you must manually add the `<component-scoped-refs>` element as described in step 7 of this topic.

If the Web service is implemented by an EJB in the same module as the client, you can add the `<port-component-link>` element as described in step 8 of this topic, to optimize access to the service.

The following steps explain how to set the fields if you choose to.

You can also configure the `webservicesclient.xml` deployment descriptor with the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>). See *Configuring the webservicesclient.xml deployment descriptor with the Assembly Toolkit*.

To configure the `webservicesclient.xml` deployment descriptor with command-line tools:

### Steps for this task

1. Open the `webservicesclient.xml` deployment descriptor in a text editor.
2. Locate the `<description>` tag in the `<service-ref>` element.
  - a. Fill in this element with a descriptive name for the service the client accesses.
3. Locate the `<service-ref-name>` tag in the `<service-ref>` element.
  - a. Fill in this element with the name the Java Naming Directory Interface (JNDI) uses to locate the service.

The JNDI lookup string for this service is

```
java:comp/env/service-ref-name
```

by convention, `service-ref-name` always begins with `service/`.
4. Locate the `<service-interface>` tag in the `<service-ref>` element.
  - a. Set this element to the class name, including package, of the generated Java interface that is the Service Interface for this Web service.
5. In the `<service-ref>` element, locate the `<wsdl-file>` tag.
  - a. Set this element to the WSDL file name used by this client relative to the root of the module.
6. Locate the `<jaxrpc-mapping-file>` tag in the `<service-ref>` element.
  - a. Set this element to the file name of the generated Java mapping file relative to the root of the module.

7. In the <port-component-ref> element, locate the <service-endpoint-interface> tag.
  - a. Set this element to the class name, including package, of the Java interface that is the Service Endpoint Interface for this Web service.
8. (Optional) Repeat the <service-ref> entry as many times as needed, filling in the elements listed in the previous steps, if your client uses more than one Web service.
9. (Optional) Define the <service-ref> element within the <component-scoped-refs> element if the service reference is being defined within an EJB JAR file, so that the service reference can be associated with the EJB.
 

This association is made by using `ejb-name` from the `ejb-jar.xml` file as the value of the `component-name` in the `component-scoped-refs` element. For more information, see section 7.2.2 of the Web services for J2EE specification available through Web services: Resources for learning.
10. (Optional) Add the <port-component-link> to optimize access to the service if the service reference is defined in an EJB JAR file that also contains the EJB that implements the service.
 

The <port-component-link> value is set to the <port-component-name> defined in `webservices.xml` in the same module.

### Results

The `webservicesclient.xml` deployment descriptor is configured.

### Usage scenario

For the AddressBook samples, the following edits are made to `webservicesclient.xml`:

1. Set the <description> element.
2. The default values are used for the remaining elements.

The resulting file contains :

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE webservicesclient PUBLIC "-//IBM Corporation, Inc.//DTD J2EE Web services client 1.0//EN"
"http://www.ibm.com/webservices/dtd/j2ee_web_services_client_1_0.dtd">
<webservicesclient>
  <service-ref>
    <description>WSDL Service AddressBookService </description>WSDL Service AddressBookService</description>
    <service-ref-name>service/AddressBookService
      <service-ref-name>service/AddressBookService</service-ref-name>
    <service-interface>com.ibm.websphere.samples.webservices.addr.AddressBookService</service-interface>
    <wsdl-file>META-INF/wsdl/AddressBook.wsdl</wsdl-file>
    <jaxrpc-mapping-file>META-INF/AddressBook_mapping.xml</jaxrpc-mapping-file>
    <port-component-ref>
      <service-endpoint-interface>com.ibm.websphere.samples.webservices.addr.AddressBook</
        service-endpoint-interface>
    </port-component-ref>
  </service-ref>
</webservicesclient>
```

### What to do next

Assemble a Web services-enabled client.

## **Configuring the webservicesclient.xml deployment descriptor with the Assembly Toolkit**

### Before you begin

You should have an enterprise JavaBean (EJB) Java archive (JAR) file, Web archive (WAR) file or an application client file that you can import into the Assembly Toolkit.

This topic explains how to configure the `webservicesclient.xml` deployment descriptor with the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>). For more information about completing tasks with the Assembly Toolkit, click **Help > Help** in the Assembly Toolkit graphical user interface (GUI).

You can also configure the `webservicesclient.xml` deployment descriptor with command-line tools. See [Configuring the webservicesclient.xml deployment descriptor with command-line tools](#).

To configure the `webservicesclient.xml` deployment descriptor with the Assembly Toolkit:

#### Steps for this task

1. Start the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>).
2. Click **File > Import** to import the EJB JAR file, WAR file or application client file into the Assembly Toolkit.
3. Right-click on the `webservicesclient.xml` file in the Navigator pane.  
You might have to switch to the Java perspective for the `webservicesclient.xml` file to be visible in the Navigator pane.
4. Select **Open**.  
**The Web Services Client** editor opens.
5. Expand the **Service references** section.
6. Select the *service\_reference* that you want to configure.
7. Expand the **Service reference overview** section.
8. Type the name of the service for which the client accesses in the **Description** field.
9. Expand the **Service reference implementation details** section.
  - a. Type the name that the Java Naming Directory Interface (JNDI) uses to locate the service in the **Service references name** field.  
The JNDI lookup string for this service is `java:comp/env/service-ref-name`. By convention, the service reference name always begins with `service/`.
  - b. Type the class name, including package, of the generated Java interface that is the Service Interface for this Web service in the **Service interface name** field.
  - c. Type the WSDL file name used by the client, relative to the root of the module, in the **WSDL file** field.
  - d. Type the file name of the Java mapping file, relative to the root of the module, in the **JAX RPC mapping file** field.

#### Results

The `webservicesclient.xml` deployment descriptor is configured using the Assembly Toolkit.

#### What to do next

## Java2WSDL command

The **Java2WSDL** command tool maps a Java class to a Web Services Description Language (WSDL) file by following the Java API for XML-based remote procedure call (JAX-RPC) specification. The **Java2WSDL** command accepts a Java class as input and produces a WSDL file representing the input class. If there is an existing file at the output location, it is overwritten. The WSDL file generated by the **Java2WSDL** command contains WSDL and XML schema constructs that are automatically derived from the input class. You can override these default values with command-line arguments.

The WSDL file generated by the **Java2WSDL** command can contain unexpected elements. Review Mapping between Java, WSDL and XML and the JAX-RPC specification available through Web services: Resources for learning, for more information on the transformations performed. You can create WSDL files that cannot be compiled when regenerated into Java code using the **WSDL2Java** command because the JAX-RPC mapping from Java to WSDL is not reversible back to the original Java code. Inspect and modify the WSDL file if you encounter this problem.

### Command line syntax and arguments

The command line syntax is:

```
Java2WSDL [argument...] class
```

The following command-line arguments are supported:

#### Required arguments

- ***class***

Represents the fully qualified name of one of the following Java classes:

- Stateless session EJB remote interface that extends the `javax.ejb.EJBObject` class
- Service Endpoint Interface that extends the `java.rmi.Remote` class
- Java bean

The **Java2WSDL** command locates the class in CLASSPATH.

#### Important arguments

- **-bindingName *name***

Specifies the name to use for the binding element. If not specified, the binding name is the `portTypeName`.

- **-help**

Displays the help message.

- **-helpX**

Displays the help message for extended options.

- **HelpXoptions**

- **-debug**

Displays debug messages.

- **-outputImpl *impl-wsdl***

Specifies if you want an interface and implementation WSDL file emitted.

- **-locationImport *location-uri***

Specifies the location of the interface WSDL file if you use the `-outputImpl` argument specified.



- **-MIMEStyle**  
Specifies a style representing Multipurpose Internet Mail Extensions (MIME) information. Valid arguments are:
  - **Axis**
  - **WSDL11** (default)
- **-soapAction**  
Valid arguments are:
  - **DEFAULT**  
Sets the soapAction field according to deployment information.
  - **NONE**  
Sets the soapAction field to "".
  - **OPERATION**  
Sets the soapAction field to the operation name.
- **-stopClasses *parent* [, *parent*]**  
If the **-all** argument is specified, the **Java2WSDL** command searches inherited classes and interfaces to construct the list of methods for WSDL file operations. The **Java2WSDL** command searches inherited classes and interfaces when generating extended complexTypes. The search stops when a class or interface is found within a package that begins with **java** or **javax**. The **-stopClasses** argument can be used to define additional classes that cause the search to stop.
- **-namespaceImpl *namespace***  
Specifies the target namespace for the implementation WSDL if **-outputImpl** specified.
- **-voidReturn**  
Valid arguments are:
  - **ONEWAY**  
Methods with void returns are one-way. This is the default for JMS transport.
  - **TWOWAY**  
Methods with void returns are two-way. This the default for HTTP transport.
- **-wrapped *boolean***  
Specifies if the WSDL file should be generated according to wrapped rules. This is only valid if use is literal. The option defaults to true.
- **-extraClasses *classes***  
Specifies other classes that should be represented in the WSDL file.
- **-input *wsdl-uri***  
Specifies the input WSDL file used to build an output WSDL file. Information from an existing WSDL file, whose name is specified in this option, is used with the input Java class to generate the desired output.
- **-implClass *impl-class***  
The **Java2WSDL** command uses method parameter names to construct the WSDL file message part names. The command automatically obtains the message names from the debug information in the class. If the class is compiled without debug information, or if the class is an interface, the method parameter names are not available. In this case, you can use the **-implClass** argument to provide an alternative class from which to obtain method parameter names. The

impl-class does not need to implement the class if the class is an interface, but it must implement the same methods as class.

- **-location *location***  
Provides the location or Uniform Resource Locator (URL) of the service. Typically, this value fills automatically when the Web service deploys. Use this argument to specify the location if you want to generate a WSDL file containing a location URL without deploying. A warning displays to remind you that the generated WSDL file should not be published if the final location is not yet been determined. The name after the last slash or backslash is the name of the service port, unless the name is overridden by the `-serviceName` argument. The service port address location attribute is assigned the specified value.
- **-namespace *targetNamespace***  
Indicates the target namespace for the WSDL file being generated. See Mapping between Java, WSDL and XML for the algorithm used to obtain the default namespace.
- **-output *wsdl-uri***  
Indicates the path and file name of the output WSDL file. If not specified, the default file, `class.wsdl`, is written into the current directory.
- **-PkgtoNS *package namespace***  
Specifies the mapping of a Java package to a namespace. If there is a package without a namespace, the **Java2WSDL** command generates a namespace name. This argument can be repeated to specify mappings for multiple packages.
- **-portTypeName *name***  
Specifies the name to use for the portType element. If not specified, the class name is used.
- **-serviceName *name***  
Specifies the name of the service element.
- **-serviceName *name***  
Specifies the name of the service. If not specified, the service name is derived from the `-location` argument.
- **-style *RPC | DOCUMENT***  
Specifies the WSDL style to use in the generated WSDL file. For more information about styles, see Mapping between Java, WSDL and XML. This argument is used with the `-use` argument.  
If `RPC` is specified with `-use ENCODED`, or omitting `use`, a `style=rpc/use=encoded` WSDL file is generated. If `RPC` is specified with `-use LITERAL`, a `style=rpc/use=literal` WSDL file is generated. If `DOCUMENT` is specified with `-use LITERAL` or omitting `use`, a `style=document/use=literal` WSDL file is generated.
- **-transport *http | jms***  
Generates Simple Object Access Protocol (SOAP) bindings for either Hyper Text Transport Protocol (HTTP) (default) or Java Messaging Service (JMS). If `jms` is specified, the characters "jms" are appended to the WSDL file name to prevent overwriting an existing WSDL file for another transport. The transport option can only be specified once.
- **-use *LITERAL | ENCODED***  
Specifies which style and use combinations are generated into the WSDL file when used with the `-style` argument. The combinations are `rpc` and `encoded`, `rpc` and `literal`, or `doc` and `literal`. For more information, see the Mapping between Java, WSDL and XML.

- **-verbose**  
Displays verbose messages.

## Developing a Web service using a stateless session enterprise bean

### Before you begin

To use an enterprise bean as the basis for a Web service implementation, follow these requirements:

- The enterprise bean must be a stateless session bean.
- Web service method parameters must be serializable and cannot be object references.
- Web service method parameters must be one of the supported Java API for XML-based remote procedure call (JAX-RPC) types.

These requirements are documented in the JAX-RPC specification available through Web services: Resources for learning

Create the artifacts that enable the enterprise bean to be a Web service and assemble the artifacts into the enterprise application as follows:

### Steps for this task

1. Access an existing Java archive (JAR) file to be used as a Web service.  
Make sure that the enterprise bean meets the requirements.
2. Develop a Service Endpoint Interface.  
The Service Endpoint Interface defines which enterprise bean methods should be made available as a Web service.
3. Develop a Web Services Description Language (WSDL) file.
4. Develop Web service deployment descriptor templates from the WSDL file.
  - a. Configure the webservices.xml deployment descriptor.
  - b. Configure the ibm-webservices-bnd.xmi deployment descriptor.
5. Assemble a Web services-enabled JAR file.
6. Assemble a Web services-enabled enterprise archive (EAR) file.
7. Enable the EAR file.  
If the EAR file contains EJB modules, it must have the Web services endpoint Web archive (WAR) file added with the endptEnabler tool before it is deployed.
8. Deploy the EAR file into WebSphere Application Server.

### Results

A Web service from a stateless session enterprise bean.

## Developing a new Web service with an existing Web Services Description Language file using a Java bean

### Before you begin

Locate the Web Services Description Language (WSDL) file that defines the Web service to be implemented. You can develop a WSDL or obtain one from an existing Web service through Email, downloading or a Uniform Resource Locator (URL).

To develop a new Web service with an existing WSDL file using a Java bean:

#### Steps for this task

1. Develop implementation templates and bindings from a WSDL file.
2. Complete the Java bean implementation.
3. Configure the `webservices.xml` deployment descriptor.
4. Configure the `ibm-webservices-bnd.xmi` deployment descriptor.
5. Assemble a Web services-enabled Web archive (WAR) file when starting from a WSDL file.
6. Assemble a Web services-enabled enterprise archive (EAR) file.
7. Deploy the EAR file into WebSphere Application Server.

### **Developing implementation templates and bindings from a Web Services Description Language file**

#### Before you begin

To develop the implementation templates and bindings from a Web Services Description (WSDL) file, you must obtain the Uniform Resource Locator (URL) of the WSDL file to use.

If it is a local file and you are running the Windows platform, the URL looks like this: `file:drive:\path\file_name.wsdl`. If you are using the UNIX platform, the URL looks like this: `file:/path/file_name.wsdl`. You can also specify local files using the absolute or relative file system path.

Implementation templates are generated using the `-role develop-server` option in combination with the `-container` option of the `WSDL2Java` command. This option takes the following parameters:

- **-container EJB**  
Generates templates for an EJB implementation in an EJB module.
- **-container Web**  
Generates templates for a Java bean implementation in a Web module.

Templates are generated for an EJB implementation for the following:

- EJB
- EJB remote interface
- EJB Home

The `WSDL2Java` command also generates bindings and deployment descriptors.

To develop implementation templates and bindings from a WSDL file:

#### Steps for this task

1. Run the `WSDL2Java -verbose -role develop-server -container type wsdlURL` command.  
Where: `type` is EJB for an enterprise JavaBean (EJB)-based implementation or Web for a Java bean-based implementation.

Since the **verbose** option is specified, a list of all generated files is displayed when the command runs.

## Results

Templates for the implementation and deployment descriptors required to implement a Web service, as well as bindings files. These templates are partially filled with information from the WSDL file.

## Usage scenario

The following example uses an enterprise bean named `AddressBook` and a WSDL file named `AddressBook.wsdl`. After generating the template files from the **WSDL2Java -verbose -role develop-server -container EJB AddressBook.wsdl** command, the following files are generated:

```
Parsing XML file: file:e:/example/app/topdown/step1/AddressBook.wsdl
WSWS3185I: Info: Parsing XML file: AddressBookW2JB.wsdl
WSWS3282I: Info: Generating addr\Address.java.
WSWS3282I: Info: Generating addr\Phone.java.
WSWS3282I: Info: Generating addr\StateType.java.
WSWS3282I: Info: Generating addr\AddressBook.java.
WSWS3282I: Info: Generating addr\AddressBookSoapBindingImpl.java.
WSWS3282I: Info: Generating addr\AddressBook_RI.java.
WSWS3282I: Info: Generating addr\AddressBookHome.java.
WSWS3282I: Info: Generating META-INF\webservices.xml.
WSWS3282I: Info: Generating META-INF\ibm-webservices-bnd.xmi.
WSWS3282I: Info: Generating META-INF\AddressBookW2JB_mapping.xml.
WSWS3282I: Info: Generating META-INF\ibm-webservices-ext.xmi.
```

If you are using a Java bean, you can generate the files with the **WSDL2Java -verbose -role develop-server -container Web AddressBook.wsdl** command.

## What to do next

Complete the EJB implementation.

Complete the Java bean implementation.

## **Completing the enterprise bean implementation**

### Before you begin

Develop implementation templates and bindings from a Web Services Description Language (WSDL) file.

To complete the enterprise bean implementation:

### Steps for this task

1. Inspect the enterprise JavaBean (EJB) remote interface template, `portType_RI.java`.  
If necessary, modify the template. `portType` is the name of the `<wsdl:portType>` element in the WSDL file.
2. Inspect the EJB home interface template, `portTypeHome.java`.  
If necessary, modify the template.
3. Edit the EJB implementation template, `bindingImpl.java`.  
`binding` is the name of the `<wsdl:binding>` element in the WSDL file.
  - a. Complete the implementation of the methods in the template.

- b. (Optional) Make changes if necessary.
  - c. (Optional) Change the class name if the binding name is not acceptable.
4. Compile all the Java classes.
  5. Assemble an EJB Java archive (JAR) file.  
Assemble all the Java classes into an EJB JAR file using the typical EJB assembly tools. Include all of the classes generated from running the **WSDL2Java** command tool when developing implementation templates and bindings from a WSDL file.

### Results

An EJB JAR file containing an EJB and supporting classes created from a WSDL file.

### What to do next

Run the **WSDL2Java** command tool generated templates for the Web services deployment descriptors. Configure the `webservices.xml` deployment descriptor.

## **Completing the Java bean implementation**

### Before you begin

Develop implementation templates and bindings from a Web Services Description Language (WSDL) file.

### Steps for this task

1. Edit the enterprise JavaBean (EJB) implementation template, `bindingImpl.java`.  
`binding` is the name of the `<wsdl:binding>` element in the WSDL file.
  - a. Complete the implementation of the methods in the template.
  - b. (Optional) Make changes if necessary.
  - c. (Optional) Change the class name if the binding name is not acceptable.
2. Compile all the Java classes.
3. Assemble a Web archive (WAR) file.  
Assemble all the Java classes into a WAR file using typical Web module assembly tools. Include all of the classes generated from running the **WSDL2Java** command tool when developing implementation templates and bindings from a WSDL file.

### Results

A Java archive (JAR) file containing a Java bean and supported classes created from the WSDL file.

### What to do next

Run the **WSDL2Java** command and generate templates for the Web services deployment descriptors. Configure the `webservices.xml` deployment descriptor

## **WSDL2Java command**

The **WSDL2Java** command tool creates Java classes and deployment descriptor templates from a Web Services Description Language (WSDL) file using the Java API for XML-based remote procedure call (JAX-RPC) 1.0 specification. See Mapping between Java, WSDL and XML for more information.

## Classes and files generated

The following kinds of classes and files are generated:

- **For each portType in the WSDL document (<wsdl:portType> element tag) use a:**
  - Service Endpoint Interface
- **For each service in the WSDL document (<wsdl:service> element tag) use a:**
  - Service Interface when the -role develop-client argument is specified.
  - ServiceLocator when the -role deploy-client argument is specified.  
This class is a WebSphere product-specific implementation of the service interface, and is not used directly.
  - webservices.xml deployment descriptor template when the -role develop-server argument is specified
  - ibm-webservices-bnd.xmi deployment descriptor template when the -role develop-server argument is specified.
  - ibm-webservices-ext.xmi deployment descriptor template when the -role develop-server argument is specified.
  - *wsdlfile\_mapping.xml* JAX-RPC mapping file when the -role develop-client or -role develop-server is specified.
  - webservicesclient.xml deployment descriptor template when the -role develop-client argument is specified.
  - ibm-webservicesclient-bnd.xmi deployment descriptor template when the -role develop-client argument is specified.
  - ibm-webservicesclient-ext.xmi deployment descriptor template when the -role develop-client argument is specified.

When the role is a server role, the container argument specifies which J2EE container the implementation uses. When the -role develop-server -container ejb arguments are specified, the webservices.xml, ibm-webservices-bnd.xmi, ibm-webservicesclient-ext.xmi and the mapping file are generated into the META-INF subdirectory. When the -role develop-server -container web arguments are specified, the files are generated into the WEB-INF directory.

- **For each binding in the WSDL file (<wsdl:binding> element tag):**
  - A stub generates that implements the Service Endpoint Interface(deploy-client role)
  - An implementation template for an enterprise bean and templates for the EJB remote interface and home generate when the -role develop-server and -container-ejb arguments are specified.
  - An implementation template for the Java bean generates when the -role develop-server and -container-web arguments are specified.
- **Other classes and files:**
  - A Java bean representing the structure of the type when the -role develop-server or -role develop-client arguments are specified for each complexType or simpleType.
  - Three classes, \*\_Ser.java, \*\_Deser.java, and \*\_Helper.java, generate for each complexType to assist in converting the bean to Simple Object Access Protocol (SOAP) and back when the -role deploy-server or -role deploy-client argument is specified.
  - A \*Holder.java class generates when the -role develop-server or -role develop-client arguments are specified for each out and inout parameter.

## Command line syntax

The command line syntax is:

```
WSDL2Java [arguments] WSDL-URI
```

## Required arguments

- **WSDL-URI**

Specifies the location of the input WSDL document using a Universal Resource Identifier (URI). You can also use a regular file path if the WSDL file is on the local file system.

## Important arguments

- **-container** *j2ee-container*

Indicates the J2EE container to be used. Valid arguments are:

- **client**

Indicates client container.

- **ejb**

Indicates enterprise JavaBean (EJB) container.

- **none**

Indicates no container.

- **web**

Indicates Web container.

If client is role, the default argument is **none**. If server is role, the container must be **ejb** or **web**. The same container option must be used for both development and deployment.

- **-deployScope** *argument*

Indicates how to deploy the server implementation. Valid arguments are:

- **Application**

Uses one instance of the implementation class for all requests.

- **Request**

Creates a new instance of the implementation class for each request.

- **Session**

Creates a new instance of the implementation class for each session.

- **-genResolver**

Generates an `absolute-import` resolver class. The purpose of this class is to record the contents of the imported WSDL files used by the WSDL URI. This class is used by the runtime. It can also be used for future **WSDL2Java** command runs. This is desirable when the imported WSDL files are remote and can be inaccessible or slow to access. It also eliminates the possibility that a remote WSDL file might have different contents at run time than it did at development time. The generated class is named `_AbsoluteImportResolver.java`. You should compile and package this class with the other Java classes generated by the **WSDL2Java** command.

- **-help**

Displays a help message and exits.

- **-helpX**

Displays a help message for extended options and exits. The options are:

- **-all**



- Generates Java files for all types, even those that are not referenced.
- **-debug**  
Prints debugging information.
- **-fileNStoPkg** *filename*  
Specifies the file of namespace to package mappings. The default is `NStoPKG.properties`.
- **-genJava** *argument*  
Generates Java files. Valid arguments are:
  - **IfNotExists**, default
  - **Overwrite**
  - **No**
- **-genXML** *argument*  
Generates the `.xml` and `.xmi` files. Valid arguments are:
  - **IfNotExists**, default
  - **Overwrite**
  - **No**
- **-password** *password*  
Specifies the login user password to access the WSDL URI.
- **-testCase**  
Generates the template for a JUnit test case for testing a Web service.
- **-user** *id*  
Specifies the login user name to access the WSDL URI.
- **-inputMappingFile** mapping file  
Specifies the file name of the Java to WSDL mapping file.
- **-NStoPkg** *namespace=package*  
By default, package names are automatically derived from the namespace strings in the WSDL file. For example, if the namespace is of the form `http://x.y.com` or `urn:x.y.com`, the corresponding package is `com.y.x`.  
You can provide your own mapping by using the **-NStoPkg** argument, which you can repeat as often as necessary, once for each unique namespace mapping. For example, if there is a namespace in the WSDL file called `urn:AddressFetcher2`, and you want files generated from the objects in this namespace to reside in the package `samples.addr`, provide the `-NStoPkg urn:AddressFetcher2=samples.addr` argument to the **WSDL2Java** command.
- **-output** directory  
Sets the root directory for emitted files.
- **-role** *j2ee role*  
Specifies the J2EE development role that identifies which files to generate. Valid arguments are:
  - **client**  
Combination of `develop-client` and `deploy-client`.
  - **deploy-client**  
Generates binding files for client deployment.
  - **deploy-server**  
Generates binding files for server deployment.
  - **develop-client** (default)  
Generates files for client development.

- **develop-server**  
Generates files for server development.
- **server**  
Combination of develop-server and deploy-server.
- **-timeout** *seconds*  
Specifies how long the **WSDL2Java** command should wait, in seconds, for the WSDL-URI to respond before giving up. The default is 45 seconds, -1 disables the timeout.
- **-useResolver** *resolver-class*  
Specifies an absolute-import resolver class to use during parsing. This class must have been created during a previous execution of the **WSDL2Java** command using the **-genResolver** option. The class must be available in CLASSPATH.
- **-verbose**  
Displays processing information, including the names of the generated files.

### Troubleshooting

The default XML namespace to Java package mapping does not take the context root into account. If two namespaces are the same up to the first slash, they are mapped to the same Java package. For example, the XML namespaces `http://www.ibm.com/foo` and `http://www.ibm.com/bar` both map to the Java package `com.ibm.www`. Use the **-NStoPkg** option to specify the package for the fully qualified namespace.

## Developing a new Web service with an existing Web Services Description Language file using a stateless session enterprise bean

### Before you begin

Locate the Web Services Description Language (WSDL) file that defines the Web service to implement. The Simple Object Access Protocol (SOAP) address URI is not required because it is updated when your new implementation is deployed.

Create the enterprise bean and artifacts that enable the enterprise bean to be a Web service and assemble those artifacts into the enterprise application as follows:

### Steps for this task

1. Develop implementation templates and bindings from a WSDL file.
2. Complete the enterprise bean implementation.
3. Configure the `webservices.xml` deployment descriptor.
4. Configure the `ibm-webservices-bnd.xmi` deployment descriptor.
5. Assemble a Web services-enabled enterprise JavaBean (EJB) Java archive (JAR) file.
6. Assemble a Web services-enabled enterprise archive (EAR) file.
7. Enable the EAR file.  
If the EAR file contains EJB modules, it must have the Web services endpoint Web archive (WAR) file added with the **endptEnabler** tool before deployment.
8. Deploy the EAR file into WebSphere Application Server.

## Results

An EJB implementation of a Web service defined in the WSDL file.

## **Web services development artifacts**

*Development artifacts* enable an enterprise bean or a Java bean module to be a Web service. To create a Web service from an enterprise bean or a Java bean module, the following files are added to the respective Java archive (JAR) or Web archive (WAR) modules at assembly time:

- **Web Services Description Language (WSDL) XML**

The WSDL XML file describes the Web service being implemented.

- **Service Endpoint Interface**

A Service Endpoint Interface is the Java interface corresponding to the Web service port type implemented. The Service Endpoint Interface is defined by the WSDL 1.1 W3C Note.

- **webservices.xml**

The `webservices.xml` file contains the J2EE Web service deployment descriptor specifying how the Web service is implemented. The `webservices.xml` file is defined in the Web services for J2EE specification available through Web services: Resources for learning

- **ibm-webservices-bnd.xmi**

This file contains WebSphere product-specific deployment information and is defined in Web services assembly properties.

- **Java API for XML-based remote procedure call (JAX-RPC) mapping file**

The JAX-RPC mapping deployment descriptor specifies how Java elements are mapped to and from WSDL file elements. The JAX-RPC mapping file is defined in the Web services assembly properties.

The following files are added to the application client module at assembly, allowing a J2EE application client to access Web services:

- **WSDL**

The WSDL file is provided by the Web service implementer.

- **Java interfaces for the Web service**

The Java interfaces are generated from the WSDL file as specified by the JAX-RPC. These bindings are the Service Endpoint Interface based on the WSDL port type, or the service interface, which is based on the WSDL service.

- **webservicesclient.xml**

The `webservicesclient.xml` file is the client-side deployment descriptor describing the services being accessed. The `webservicesclient.xml` file is defined in the Web services for J2EE specification, available through Web services: Resources for learning.

- **ibm-webservicesclient-bnd.xmi**

This file contains WebSphere product-specific deployment information such as security information. The `ibm-webservicesclient-bnd.xmi` file is defined in the Web services assembly properties.

- **Other JAX-RPC binding files** Additional JAX-RPC binding files that support the client application in mapping Simple Object Access Protocol (SOAP) to Java language are generated from WSDL by the **WSDL2Java** command tool.

# Mapping between Java, Web Services Description Language and XML

This topic contains the mappings between Java and XML technologies, including XML Schema, Web Services Description Language (WSDL) and Simple Object Access Protocol (SOAP), supported by WebSphere Application Server. Most of these mappings are specified by Java API for XML-based remote procedure call (JAX-RPC). Some mappings optional or unspecified in JAX-RPC are also supported.

There are references to the JAX-RPC specification through this topic. You can review the JAX-RPC specification through [Web services: Resources for learning](#).

## Notational conventions

The following table specifies the namespace prefixes and corresponding namespaces used.

| Namespace prefix | Namespace   |
|------------------|---|
| xsd              | <a href="http://www.w3.org/2001/XMLSchema">http://www.w3.org/2001/XMLSchema</a>                   |
| xsi              | <a href="http://www.w3.org/2001/XMLSchema-instance">http://www.w3.org/2001/XMLSchema-instance</a> |
| soapenc          | <a href="http://schemas.xmlsoap.org/soap/encoding/">http://schemas.xmlsoap.org/soap/encoding/</a> |
| wsdl             | <a href="http://schemas.xmlsoap.org/wsdl/">http://schemas.xmlsoap.org/wsdl/</a>                   |
| wsdlsoap         | <a href="http://schemas.xmlsoap.org/wsdl/soap/">http://schemas.xmlsoap.org/wsdl/soap/</a>         |
| ns               | user defined namespace  |
| apache           | <a href="http://xml.apache.org/xml-soap">http://xml.apache.org/xml-soap</a>                       |
| wasws            | <a href="http://websphere.ibm.com/webservices/">http://websphere.ibm.com/webservices/</a>         |

## Detailed mapping information

The following sections identify the supported mappings, including:

- Java-to-WSDL mapping
- WSDL-to-Java mapping
- Mapping between WSDL and SOAP messages

### Java-to-WSDL mapping

This section summarizes the Java-to-WSDL mapping rules. The Java-to-WSDL mapping rules are used by the **Java2WSDL** command tool for bottom-up processing. In bottom-up processing, an existing Java service implementation is used to create a WSDL file defining the Web service. The generated WSDL file can require additional manual editing for the following reasons:

- Not all Java classes and constructs have mappings to WSDL. For example, Java classes that do not comply with the Java bean specification rules might not map to a WSDL construct.
- Some Java classes and constructs have multiple mappings to WSDL. For example, a `java.lang.String` class can be mapped to either an `xsd:string` or `soapenc:string`. The **Java2WSDL** command chooses one of these mappings, but the WSDL file must be edited if a different mapping is desired.

- There are multiple ways to generate WSDL constructs. For example, the part element in the `wsdl:message` can be generated with a type or element attribute. The **Java2WSDL** command makes an informed choice based on the settings of the `-style` and `-use` options.
- The WSDL file describes the instance data elements sent in the SOAP message. If you want to modify the names or format used in the message, the WSDL file must be edited. For example, write a bean property value as an attribute instead of an element.
- The WSDL file requires editing if header or attachment support is desired.
- The WSDL file requires editing if a multipart WSDL, one using `wsdl:import`, is desired.

For simple services, the generated WSDL file is sufficient. For complicated services, the generated WSDL file is a good starting point.

### General issues

- **Package to namespace mapping**

The JAX-RPC does not specify the default mapping of Java package names to XML namespaces. The JAX-RPC does specify that each Java package must map to a single XML namespace, and likewise. A default mapping algorithm is provided that constructs the namespace by reversing the names of the Java package and adding an `http://` prefix. For example, a package named, `com.ibm.webservice`, is mapped to the namespace `http://webservice.ibm.com`.

The default mapping between XML namespaces and Java package names can be overridden using the `-NStoPkg` and `-PkgtoNS` options of the **WSDL2Java** and **Java2WSDL** commands.

- **Identifier mapping**

Java identifiers are mapped directly to WSDL file and XML identifiers.

Java bean property names are mapped to the WSDL file and XML identifiers. For example, a Java bean, with `getInfo` and `setInfo` methods, maps to an XML construct with the name, `info`.

The Service Endpoint Interface method parameter names, if available, are mapped directly to the XML identifiers. See the **WSDL2Java** command `-implClass` option for more details.

- **WSDL construction summary**

The following table summarizes the mapping from a Java construct to the related WSDL and XML construct.

| Java construct             |   |
|----------------------------|---|
| Service Endpoint Interface | <code>wsdl:portType</code>  |
| Method                     | <code>wsdl:operation</code>   |
| Parameters                 | <code>wsdl:input</code> , <code>wsdl:message</code> , <code>wsdl:part (1)</code>  |
| Return                     | <code>wsdl:output</code> , <code>wsdl:message</code> , <code>wsdl:part (1)</code> |
| Throws                     | <code>wsdl:fault</code> , <code>wsdl:message</code> , <code>wsdl:part (1)</code>  |
| Primitive types            | <code>xsd</code> and <code>soapenc</code> simple types                            |
| Java beans                 | <code>xsd:complexType</code>  |
| Java bean properties       | Nested <code>xsd:elements</code> of <code>xsd:complexType</code>                  |
| Arrays                     | JAX-RPC defined array <code>xsd:complexType</code>                                |
| User defined exceptions    | <code>xsd:complexType</code>  |

**Note:** The generated WSDL file is affected by the `-style` and `-use` options. A `wsdl:binding` that conforms to the generated `wsdl:portType` is generated. The style and use constructs of the `wsdl:binding` are determined from the `-style` and `-use` options. A `wsdl:service` containing a port that references the generated `wsdl:binding` is generated. The names and values of the `wsdl:service` are controlled by the **Java2WSDL** command options.

- **Style and use**

Use the `-style` and `-use` options to generate different kinds of WSDL files. The four supported combinations are:

- `-style RPC -use ENCODED`
- `-style DOCUMENT -use LITERAL`
- `-style RPC -use LITERAL`
- `-style DOCUMENT -use LITERAL -wrapped false`

The `-use LITERAL` option affects the generated WSDL file in the following ways:

- The `soap:body` elements in the `wsdl:binding` are specified as `use="literal"`.
- The `soap:fault` elements in the `wsdl:binding` are specified as `use="literal"`.
- The soap encoded types are not used.
- The soap encoded array style is not used. The `maxOccurs` attribute is used to indicate arrays.

The `-use ENCODED` option affects the generated WSDL file in the following ways:

- The `soap:body` elements in the `wsdl:binding` are specified as `use="encoded"` and the `encodingStyle` is set.
- The `soap:fault` elements in the `wsdl:binding` are specified as `use="encoded"` and the `encodingStyle` is set.

The `-style RPC` option affects the generated WSDL file in the following ways:

- The `wsdl:part` elements use the `type` attribute to reference XML types.
- The `wsdl:binding` is specified as `style="rpc"`.

The `-style DOCUMENT -wrapped false` option affects the generated WSDL file in the following ways:

- The `wsdl:part` elements use the `type` attribute to reference simple types. The `element` attribute is used to reference the root `xsd:elements` for everything that is not a simple type.
- The `wsdl:binding` is specified as `style="document"`.

The `-style DOCUMENT -wrapped true` option generates a WSDL file that conforms to the .NET WSDL file format:

- A request `xsd:element` is generated for each method in the Service Endpoint Interface as follows:
  - The name of the `xsd:element` is the same as the name of the `wsdl:operation`.
  - The `xsd:element` contains an `xsd:sequence` that contains `xsd:elements` defining each parameter.
  - The request `wsdl:message` references the wrapper `xsd:element` using a single part:
    - The name of the part is `parameters`.
    - The `element` attribute is used to reference the wrapper `xsd:element`.
- A response `xsd:element` is generated for each method in the Service Endpoint Interface as follows:

- The name of the `xsd:element` is the same as the name of the `wsdl:operation` appended with `Response`
- The `xsd:element` contains an `xsd:sequence` that contains `xsd:elements` defining the return value.
- The request `wsdl:message` references this wrapper `xsd:element` using a single part.
  - The `element` attribute is used to reference the wrapper `xsd:element`.
- The `wsdl:binding` is specified as `style="document"`.

### Mapping of standard XML types from Java types

Some Java types map directly to standard XML types. These types do not require additional XML definitions in the `wsdl:types` section.

- **JAX-RPC Java primitive type mapping**

The following table describes the mapping from the Java primitive and standard types to XML standard types. For more information see the JAX-RPC specification.

| Java type  | XML type  |
|--|---|
| <code>int</code>   | <code>xsd:int</code>  |
| <code>long</code>  | <code>xsd:long</code>   |
| <code>short</code>                                       | <code>xsd:short</code>  |
| <code>float</code>                                       | <code>xsd:float</code>  |
| <code>double</code>                                      | <code>xsd:double</code>   |
| <code>boolean</code>                                     | <code>xsd:boolean</code>  |
| <code>byte</code>  | <code>xsd:byte</code>   |
| <code>byte[]</code>                                      | <p><code>xsd:base64Binary</code></p> <p><b>Note:</b> The default mapping for <code>byte[]</code> is <code>xsd:base64Binary</code>. The data in <code>byte[]</code> is passed over the wire as a text string encoded in the base64 format. An alternative format is <code>xsd:hexBinary</code>. To use the <code>xsd:hexBinary</code> format:</p> <ul style="list-style-type: none"> <li>- Edit the WSDL file and change <code>xsd:base64Binary</code> to <code>xsd:hexBinary</code>, or</li> <li>- Change your implementation to use the specialized <code>com.ibm.ws.webservices.engine.types.HexBinary</code> class.</li> </ul> |
| <code>java.lang.String</code>                            | <code>xsd:string</code>   |
| <code>java.math.BigInteger</code>                        | <code>xsd:integer</code>  |
| <code>java.math.BigDecimal</code>                        | <code>xsd:decimal</code>  |
| <code>java.util.Calendar</code>                          | <code>xsd:dateTime</code>   |
| <code>java.util.Date</code>                              | <code>xsd:date</code>   |
| <b>Note:</b> This mapping is not covered by the JAX-RPC. |   |
| <code>java.lang.Boolean</code>                           | <code>xsd:boolean xsi:nillable=true</code>  |
| <code>java.lang.Float</code>                             | <code>xsd:float xsi:nillable=true</code>  |

|                   |                              |
|-------------------|------------------------------|
| java.lang.Double  | xsd:double xsi:nillable=true |
| java.lang.Integer | xsd:int xsi:nillable=true    |
| java.lang.Short   | xsd:short xsi:nillable=true  |
| java.lang.Byte    | xsd:byte xsi:nillable=true   |

**Note:** The java.lang wrapper classes in the last six lines of the table map to the same XML construct as the corresponding Java primitive type. In addition, the xsi:nillable attribute is generated to indicate that such elements can be null.

- **Additional Java class mappings**

In addition to the standard JAX-RPC mapping, the following classes are mapped directly to XML types:

| Java type  | XML type               |
|--|------------------------|
| com.ibm.ws.webservices.engine.types.HexBinary  | xsd:hexBinary          |
| —  | —                      |
| javax.xml.namespace.QName  | xsd:qname              |
| —  | —                      |
| com.ibm.ws.webservices.engine.types.Token  | xsd:token              |
| com.ibm.ws.webservices.engine.types.NormalizedString                                 | xsd:normalizedString   |
| com.ibm.ws.webservices.engine.types.Name   | xsd:Name               |
| com.ibm.ws.webservices.engine.types.NCName   | xsd:NCName             |
| com.ibm.ws.webservices.engine.types.NMTOKEN  | xsd:NMTOKEN            |
| com.ibm.ws.webservices.engine.types.URI  | xsd:anyURI             |
| —  | —                      |
| com.ibm.ws.webservices.engine.types.UnsignedLong                                     | xsd:unsignedLong       |
| com.ibm.ws.webservices.engine.types.UnsignedInt                                      | xsd:unsignedInt        |
| com.ibm.ws.webservices.engine.types.UnsignedByte                                     | xsd:unsignedByte       |
| com.ibm.ws.webservices.engine.types.NonNegativeInteger                               | xsd:nonNegativeInteger |
| com.ibm.ws.webservices.engine.types.PositiveInteger                                  | xsd:positiveInteger    |
| com.ibm.ws.webservices.engine.types.NonPositiveInteger                               | xsd:nonPositiveInteger |
| —  | —                      |
| com.ibm.ws.webservices.engine.types.Time   | xsd:time               |
| com.ibm.ws.webservices.engine.types.YearMonth  | xsd:gYearMonth         |
| com.ibm.ws.webservices.engine.types.Year   | xsd:gYear              |
| com.ibm.ws.webservices.engine.types.Month  | xsd:gMonth             |
| com.ibm.ws.webservices.engine.types.Day  | xsd:gDay               |
| com.ibm.ws.webservices.engine.types.MonthDay   | xsd:gMonthDay          |
| com.ibm.ws.webservices.engine.types.Duration   | xsd:duration           |
| —  | —                      |
| java.util.Map  | apache:Map             |
| <b>Note:</b> Any classes that implement java.util.Map are also mapped to apache:Map. |                        |



|  |                    |
|--|--------------------|
| java.util.Collection<br><b>Note:</b> Each Java array, except byte[], and every class that implements java.util.Collection is mapped to a JAX-RPC defined soapenc:Array type. | soapenc:Array      |
| org.w3c.dom.Element  | apache:Element     |
| java.util.Vector   | apache:Vector      |
| java.awt.Image<br><b>Note:</b> Used for attachment support.  | apache:Image       |
| javax.mail.internet.MimeMultiPart<br><b>Note:</b> Used for attachment support.   | apache:Multipart   |
| javax.xml.transform.Source<br><b>Note:</b> Used for attachment support.  | apache:Source      |
| javax.activation.DataHandler<br><b>Note:</b> Used for attachment support.  | apache:DataHandler |

### Generation of wsdl:types

Java types that cannot be mapped directly to standard XML types are generated in the wsdl:types section.

- **Java arrays**

Java arrays for the -use ENCODED option, with the exception of byte[], are generated using the following format. See the JAX-RPC specification for more details. Alternative mappings can be found in Table 18.1 of the JAX-RPC specification.

**Java:**

```
Item[]
```

**Mapped to:**

```
<xsd:complexType name="ArrayOfItem">
  <xsd:complexContent>
    <xsd:restriction base="soapenc:Array">
      <xsd:attribute ref="soapenc:arrayType" wsdl:arrayType="ns:Item" />
    </xsd:restriction>
  </xsd:complexContent>
</xsd:complexType>
```

- **JAX-RPC value type and bean mapping**

A Java class that matches the Java value type or bean pattern is mapped to an xsd:complexType. In order for a Java class to be mapped to XML, follow these conditions:

- The class must have a public default constructor.
- The class must not implement, directly or indirectly, java.rmi.Remote.
- Public, nonstatic, nonfinal, nontransient fields are mapped. The class can contain other fields and methods, but they are not mapped.

- If the class follows the Java bean pattern and has public getter and setter methods, the property is mapped.

Additional mapping rules can be found in the JAX-RPC specification. The following example indicates the mapping for sample base and derived Java classes:

**Java:**

```
public abstract class Base {
    public Base() {}
    public int a; // mapped
    private int b; // mapped via setter/getter
    private int c; // not mapped
    private int[] d; // mapped via indexed setter/getter

    public int getB() { return b;} // map property b
    public void setB(int b) {this.b = b;}

    public int[] getD() { return d;} // map indexed property d
    public void setD(int[] d) {this.d = d;}
    public int getD(int index) { return d[index];}
    public void setB(int index, int value) {this.d[index] = value;}

    public void someMethod() {...} // not mapped
}

public class Derived extends Base {
    public int x; // mapped
    private int y; // not mapped
}
```

**Mapped to:**

```
<xsd:complexType name="Base" abstract="true">
  <xsd:sequence>
    <xsd:element name="a" type="xsd:int" />
    <xsd:element name="b" type="xsd:int" />
    <xsd:element name="d" minOccurs="0" maxOccurs="unbounded" type="xsd:int"/>
  </xsd:sequence>
</xsd:complexType>

<xsd:complexType name="Derived">
  <xsd:complexContent>
    <xsd:extension base="ns:Base">
      <xsd:sequence>
        <xsd:element name="x" type="xsd:int" />
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>
```

***Inheritance and abstract classes***

The example contains two optional JAX-RPC features that are supported by WebSphere Application Server:

1. An abstract class is mapped to an `xsd:complexType` with `abstract="true"`.
2. An indexed bean property (see the methods for `d` in `Base`) are mapped to a nested element specified with `maxOccurs="unbounded"`. This format is similar to an XML array, but the SOAP message is different. An XML array defines an element for the array and nested elements for each item in the array. An

element defined with `maxOccurs` indicates a series of items without the surrounding array wrapper element. Both formats are popular.

- **JAX-RPC enumeration class mapping**

Section 4.2.4 of the JAX-RPC specification defines the mapping from an XML enumeration to a Java class. Though not specifically required by the JAX-RPC, the **Java2WSDL** command performs the reverse mapping. If you have a class that has the same format as a JAX-RPC enumeration class, it is mapped to an XML enumeration.

- **Holder classes**

The JAX-RPC specification defines Holder classes in section 4.3.5. A Holder class is used to support in and out parameter passing. Every Holder class implements the `javax.xml.rpc.holders.Holder` interface. The **Java2WSDL** command maps Holder classes to the same XML type as the held type. In addition, references to Holder classes affect the generation of `wsdl:messages`.

- **Exception classes**

If a class extends the exception, `java.lang.Exception`, it is mapped to an `xsd:complexType` similar to the Java bean mapping. The getter methods of the exception are mapped as nested `xsd:elements` of the `xsd:complexType`. See section 5.5.5 of the JAX-RPC specification for more details.

**Note:** You need to generate implementation specific exception classes by invoking the **WSDL2Java** command on the resulting WSDL file.

- **Unsupported classes**

If a class cannot be mapped to an XML type, the **Java2WSDL** command issues a message and an `xsd:anyType` reference is generated in the WSDL file. In these situations, modify the Web service implementation to use the JAX-RPC compliant classes.

- **Generation of root elements**

If the **Java2WSDL** command generates an `xsd:complexType` or `xsd:simpleType` for a parameter reference, a corresponding `xsd:element` is also generated. The `xsd:element` has the same name as the `xsd:complexType/xsd:simpleType` and uses the type attribute to reference the `xsd:complexType/xsd:simpleType`. The `wsdl:message` part can use the `element` attribute or the `type` attribute to reference the `xsd:element` or type. This choice is determined by the `-style` and `-use` options.

## Generation from the interface or implementation class

The class passed to the **Java2WSDL** command represents the interface of the `wsdl:service`. The `wsdl:portType` and `wsdl:message` elements generate from this interface or implementation class.

- **Generation of the `wsdl:portType`**

The name of the `wsdl:portType` is the name of the class unless overridden by the `-portTypeName` option.

- **Generation of `wsdl:operation`**

A `wsdl:operation` generates for each public method in the interface that throws the exception, `java.rmi.RemoteException`.

- The name of the `wsdl:operation` is the name of the method.
- The `wsdl:operation` has a `parameterOrder` attribute, which defines the order of the parameters in the signature. Specifically, the `parameterOrder` lists the order of the parts of the request or response `wsdl:messages`.

- The `wsdl:operation` has a nested `wsdl:input` element that references the request `wsdl:message` using the `message` attribute.
- The `wsdl:operation` has a nested `wsdl:output` element that references the response `wsdl:message` using the `message` attribute.
- The `wsdl:operation` has a nested `wsdl:fault` element that references the default `wsdl:message` using the `message` attribute.

See sections 5.5.4 and 5.5.5 of the JAX-RPC specification for more information.

- **Generation of `wsdl:message`**

Generating the `wsdl:message` is directly related to the `-style` and `-use` options. The following is the default mapping (`-style RPC -use ENCODED`):

- A `wsdl:message` is created to represent the request.
  - A `wsdl:part` representing each parameter is added to the `wsdl:message`.
    - The `wsdl:part` has the same name as the parameter.
    - The `wsdl:part` uses the `type` attribute to locate the XML type of the parameter.
- A `wsdl:message` is created to represent the response:
  - A `wsdl:part` representing the method return is created.
    - The `wsdl:part` has the same name as the method with `Return` appended.
 

**Note:** The name of the part is not specified by the JAX-RPC and is typically not checked by SOAP engines.
    - The `wsdl:part` has the same name as the parameter.
    - The `wsdl:part` uses the `type` attribute to locate the XML type of the parameter.
    - A `wsdl:part` is created for each parameter that is a Holder.
    - The `wsdl:part` has the same name as the parameter.
    - A `wsdl:message` is created to represent the fault if the operation has a `wsdl:fault`.
    - A `wsdl:part` representing the fault is created.
    - The `wsdl:part` has the same name as the exception.
    - The `wsdl:part` uses the `type` attribute to locate the `complexType` representing the exception.

The same mapping is used as described if you use the `-style RPC` and `-use LITERAL` options. It is also valid to use the `wsdl:part` element attribute instead of the `type` attribute to reference the XML schema. If you use the `-style DOCUMENT -wrapped false` and `-use LITERAL` options, the same mapping is used as described except the `wsdl:part` element attribute is used to reference the XML schema. If the XML schema is a primitive type, like `xsd:string`, the `type` attribute is used to reference the XML type. The `-style DOCUMENT, -wrapped true` and `-use LITERAL` options result in completely different mappings for the request and response messages. For example:

- A request `xsd:element` is generated for each method in the Service Endpoint Interface.
  - The name of the `xsd:element` is the same as the name of the `wsdl:operation`.
  - The `xsd:element` contains an `xsd:sequence` that contains `xsd:elements` defining each parameter.
  - The request `wsdl:message` references the wrapper `xsd:element` using a single part.
    - The name of the part is `parameters`.

- The `element` attribute is used to reference the wrapper `xsd:element`.
- A response `xsd:element` is generated for each method in the Service Endpoint Interface.
  - The name of the `xsd:element` is the same as the name of the `wSDL:operation` appended with `Response`.
  - The `xsd:element` contains an `xsd:sequence` that contains `xsd:elements` defining the return value.
  - The request `wSDL:message` references this wrapper `xsd:element` using a single part.
    - The `element` attribute is used to reference the wrapper `xsd:element`.
- **Generation of `wSDL:binding`**  
 Generate a `wSDL:binding` with a name defined by the **Java2WSDL** **-bindingName** command.
  - The `wSDLsoap:binding` `style` attribute is set to `rpc` if you use the `-style RPC` option; otherwise it is set to `document`.
  - A `wSDL:operation` generates for each `wSDL:operation` defined in the `wSDL:portType`.
  - Each `wSDL:operation` has corresponding `wSDL:input`, `wSDL:output` and `wSDL:fault` elements.
  - The `wSDL:input`, `wSDL:output` and `wSDL:fault` elements each contain a `wSDLsoap:body` element.
  - The `wSDLsoap:body` `use` attribute is set to `literal` or `encoded` according to the **-use** argument. Set the `encodingStyle` attribute to `http://schemas.xmlsoap.org/soap/encoding/` when **use** is `encoded`.
- **Generation of the `wSDL:service`**  
 Generate a `wSDL:service` with a name defined by the **Java2WSDL** **-serviceElement** command. For example:
  - The `wSDL:service` contains a port with a name defined by the **Java2WSDL** **-servicePortName** command.
  - The port references the generated `wSDL:binding` with the `binding` attribute.
  - The port contains a `wSDLsoap:address` element with a
  - The `location` attribute is set to the value of the **Java2WSDL** **-location** command.

## WSDL-to-Java mapping

The **WSDL2Java** command tool uses the following rules to generate Java classes when developing your Web services client and server. In addition, implementation specific Java classes are generated that assist in the serialization and deserialization, and invocation of the Web service.

### General issues

#### • Mapping of namespace to package

The JAX-RPC does not specify the mapping of XML namespaces to Java package names. The JAX-RPC does specify that each Java package map to a single XML namespace, and likewise. A default mapping algorithm omits any protocol from the XML namespace and reverses the names. For example, an XML namespace `http://websphere.ibm.com` becomes a Java package with the name `com.ibm.websphere`.

The default mapping of XML namespace to Java package disregards the context-root. If two namespaces are the same up until the first slash, they map to

the same Java package. For example, the XML namespaces `http://websphere.ibm.com/foo` and `http://websphere.ibm.com/bar` map to the Java package `com.ibm.websphere`. The default mapping between XML namespaces and Java package names can be overridden using the `-NStoPkg` and `-PkgtoNS` options of **WSDL2Java** and **Java2WSDL** commands.

- **Identifier mapping**

XML names are much richer than Java identifiers. They can include characters that are not permitted in Java identifiers. See section 20 of the JAX-RPC specification for the rules to map an XML name to a Java identifier.

The mapping rules attempt to follow accepted Java coding conventions. Class names always begin with an uppercase letter. Method names begin with a lowercase letter. The **WSDL2Java** command generates metadata in the `_Helper` class so that the values are serialized or deserialized using the XML names specified in the WSDL file.

- **Java construction summary**

| WSDL and XML  | Java   |
|---|--|
| <p><code>xsd:complexType</code> (struct)</p> <p><b>Note:</b> The <code>xsd:complexType</code> can also represent a Java exception if referenced by a <code>wsdl:message</code> for a <code>wsdl:fault</code>.</p>   | <p>Java Bean Class</p> <p><b>Note:</b> The classes, <code>_Helper</code>, <code>_Ser</code>, and <code>_Deser</code>, generate for each Java bean class. These implementation classes aid serialization and deserialization.</p> |
| <code>nested xsd:element/xsd:attribute</code>   | Java bean property   |
| <code>xsd:complexType</code> (array)  | Java array   |
| <code>xsd:simpleType</code> (enumeration)   | JAX-RPC enumeration class  |
| <p><code>xsd:complexType</code> (wrapper)The method parameter signature typically is determined by the <code>wsdl:message</code>. However, if the WSDL file is a .NET wrapped style, the method parameter signature is determined by the wrapper <code>xsd:element</code></p> | <p>Service Endpoint Interface method parameter signature</p> <p><b>Note:</b> If a parameter is out or inout, a Holder class generates.</p>   |
| —   | —  |
| <p><code>wsdl:message</code>The method parameter signature typically is determined by the <code>wsdl:message</code>. However, if the WSDL file is a .NET wrapped style, the method parameter signature is determined by the wrapper <code>xsd:element</code></p>              | <p>Service Endpoint Interface method signature</p> <p><b>Note:</b> If a parameter is out or inout, a Holder class generates.</p>   |
| <code>wsdl:portType</code>  | Service Endpoint Interface   |
| <code>wsdl:operation</code>   | Service Endpoint Interface method  |
| <code>wsdl:binding</code>   | <p>Stub</p> <p><b>Note:</b> The Stub and ServiceLocator classes are implementation specific.</p>   |
| <code>wsdl:service</code>   | <p>Service Interface and ServiceLocator</p> <p><b>Note:</b> The Stub and ServiceLocator classes are implementation specific.</p>   |
| <code>wsdl:port</code>  | Port accessor method in Service Interface  |

## Mapping standard XML types

- **JAX-RPC simple XML types mapping**

The following mappings are XML types to Java types. For more information about these mappings, see section 4.2.1 of the JAX-RPC specification.

| XML type  | Java type            |
|---|----------------------|
| xsd:string  | java.lang.String     |
| xsd:integer   | java.math.BigInteger |
| xsd:int<br><b>Note:</b> If an element with this type has the xsi:nil attribute set to true, it is mapped to the Java wrapper class of the primitive type.     | int                  |
| xsd:long<br><b>Note:</b> If an element with this type has the xsi:nil attribute set to true, it is mapped to the Java wrapper class of the primitive type.    | long                 |
| xsd:short<br><b>Note:</b> If an element with this type has the xsi:nil attribute set to true, it is mapped to the Java wrapper class of the primitive type.   | short                |
| xsd:decimal   | java.math.BigDecimal |
| xsd:float<br><b>Note:</b> If an element with this type has the xsi:nil attribute set to true, it is mapped to the Java wrapper class of the primitive type.   | float                |
| xsd:double<br><b>Note:</b> If an element with this type has the xsi:nil attribute set to true, it is mapped to the Java wrapper class of the primitive type.  | double               |
| xsd:boolean<br><b>Note:</b> If an element with this type has the xsi:nil attribute set to true, it is mapped to the Java wrapper class of the primitive type. | boolean              |
| xsd:byte<br><b>Note:</b> If an element with this type has the xsi:nil attribute set to true, it is mapped to the Java wrapper class of the primitive type.    | byte                 |
| xsd:dateTime  | java.util.Calendar   |
| xsd:date<br><b>Note:</b> This mapping is not supported by the JAX-RPC.  | java.util.Date       |
| xsd:base64Binary  | byte[]               |

|  |                      |
|--|----------------------|
| xsd:hexBinary  | byte[]               |
| —  | —                    |
| soapenc:base64   | byte[]               |
| soapenc:base64Binary                                       | byte[]               |
| soapenc:string   | java.lang.String     |
| soapenc:boolean  | java.lang.Boolean    |
| soapenc:float  | java.lang.Float      |
| soapenc:double   | java.lang.Double     |
| soapenc:decimal  | java.math.BigDecimal |
| soapenc:int  | java.lang.Integer    |
| soapenc:integer  | java.math.BigInteger |
| <b>Note:</b> This mapping is not supported by the JAX-RPC. |                      |
| soapenc:short  | java.lang.Short      |
| soapenc:long   | java.lang.Long       |
| <b>Note:</b> This mapping is not supported by the JAX-RPC. |                      |
| soapenc:byte   | java.lang.Byte       |

- **JAX-RPC optional simple XML type mapping**

The **WSDL2Java** command supports the following JAX-RPC optional simple XML types.

| XML type               | Java type  |
|------------------------|--|
| xsd:qname              | javax.xml.namespace.QName                              |
| xsd:time               | com.ibm.ws.webservices.engine.types.Time               |
| xsd:gYearMonth         | com.ibm.ws.webservices.engine.types.YearMonth          |
| xsd:gYear              | com.ibm.ws.webservices.engine.types.Year               |
| xsd:gMonth             | com.ibm.ws.webservices.engine.types.Month              |
| xsd:gDay               | com.ibm.ws.webservices.engine.types.Day                |
| xsd:gMonthDay          | com.ibm.ws.webservices.engine.types.MonthDay           |
| xsd:token              | com.ibm.ws.webservices.engine.types.Token              |
| xsd:normalizedString   | com.ibm.ws.webservices.engine.types.NormalizedString   |
| xsd:unsignedLong       | com.ibm.ws.webservices.engine.types.UnsignedLong       |
| xsd:unsignedInt        | com.ibm.ws.webservices.engine.types.UnsignedInt        |
| xsd:unsignedShort      | com.ibm.ws.webservices.engine.types.UnsignedShort      |
| xsd:unsignedByte       | com.ibm.ws.webservices.engine.types.UnsignedByte       |
| xsd:nonNegativeInteger | com.ibm.ws.webservices.engine.types.NonNegativeInteger |
| xsd:negativeInteger    | com.ibm.ws.webservices.engine.types.NegativeInteger    |
| xsd:positiveInteger    | com.ibm.ws.webservices.engine.types.PositiveInteger    |
| xsd:nonPositiveInteger | com.ibm.ws.webservices.engine.types.NonPositiveInteger |
| xsd:Name               | com.ibm.ws.webservices.engine.types.Name               |
| xsd:NCName             | com.ibm.ws.webservices.engine.types.NCName             |



|              |  |
|--------------|--|
| xsd:NMTOKEN  | com.ibm.ws.webservices.engine.types.NMTOKEN  |
| xsd:duration | com.ibm.ws.webservices.engine.types.Duration |
| xsd:anyURI   | com.ibm.ws.webservices.engine.types.URI      |

- **JAX-RPC xsd:anyType mapping**

The **WSDL2Java** command maps an `xsd:anyType` to a `java.lang.Object`. This is an optional feature of the JAX-RPC specification. The `xsd:anyType` can be used to store any XML type other than the XML primitive type. An `xsd:anyType` is always serialized along with an `xsi:type` that specifies the actual type.

- **Additional supported mappings**

The following mappings are also supported by the **WSDL2Java** command. These mappings are not defined by the JAX-RPC specification.

| XML type   | Java type                                  |
|--|--|
| apache:PlainText<br><b>Note:</b> For MIME attachments.   | java.lang.String                           |
| apache:Map   | java.util.Map                              |
| apache:Element   | org.w3c.dom.Element                        |
| wasws:SOAPElement  | com.ibm.ws.webservices.xmlsoap.SOAPElement |
| apache:Vector  | java.util.Vector                           |
| apache:Image<br><b>Note:</b> For MIME attachments.       | java.awt.Image                             |
| apache:Multipart<br><b>Note:</b> For MIME attachments.   | javax.mail.internet.MimeMultipart          |
| apache:Source<br><b>Note:</b> For MIME attachments.      | javax.xml.transform.Source                 |
| apache:octetStream<br><b>Note:</b> For MIME attachments. | javax.activation.DataHandler               |
| apache:DataHandler<br><b>Note:</b> For MIME attachments. | javax.activation.DataHandler               |

### Mapping XML defined in the `wsdl:types` section

The **WSDL2Java** command generates Java types for the XML schema constructs defined in the `wsdl:types` section. The XML schema language is broader than the required or optional subset defined by the JAX-RPC specification. The **WSDL2Java** command supports all required mappings and most optional mappings. In addition, the command supports some XML schema mappings that are outside the JAX-RPC specification. In general, the **WSDL2Java** command ignores constructs that it does not support. For example, the **WSDL2Java** command does not support the `default` attribute. If an `xsd:element` is defined with the `default` attribute, the `default` attribute is ignored. In some cases it maps unsupported constructs to `wasws:SOAPElement`.

- **Mapping of `xsd:complexType` to Java bean**

The most common mapping is from an `xsd:complexType` to a Java bean class.

– *Standard Java bean mapping*

The standard Java bean mapping is defined in section 4.2.3 of the JAX-RPC specification. The `xsd:complexType` defines the type. The nested `xsd:elements` within the `xsd:sequence` or `xsd:all` groups are mapped to Java bean properties. For example:

**XML:**

```
<xsd:complexType name="Sample">
  <xsd:sequence>
    <xsd:element name="a" type="xsd:string"/>
    <xsd:element name="b" maxOccurs="unbounded" type="xsd:string"/>
  </xsd:sequence>
</xsd:complexType>
```

**Java:**

```
public class Sample {
  // ..
  public Sample() {}

  // Bean Property a
  public String getA()          {...}
  public void   setA(String value) {...}

  // Indexed Bean Property b
  public String[] getB()          {...}
  public String  getB(int index) {...}
  public void    setB(String[] values) {...}
  public void    setB(int index, String value) {...}
}
```

– *Methods equals() and hashCode()*

The generated Java bean classes contain an implementation of the `equals()` method. The generation of this method is outside the JAX-RPC specification. The `equals()` method returns true if `equals()` is true for each contained bean property. The implementation accounts for self-referencing loops. This version of the `equals()` method is typically more useful than the "identity" `equals` provided by `java.lang.Object`.

A corresponding `hashCode()` method is also generated in the Java bean class.

– *Properties and indexed properties*

In the standard Java bean mapping example, the nested `xsd:element` for property a is mapped to a Java bean property. In addition, the **WSDL2Java** command maps a nested `xsd:element` with `maxOccurs > 1` to a Java bean indexed property.

– *Attributes*

The **WSDL2Java** command also supports the `xsd:attribute` element, as shown in the following example.

Attribute a is mapped as a Java bean property, which is exactly the same mapping as a nested `xsd:element`. Implementation specific metadata is generated in the `Sample2_Helper` class to ensure that property a is serialized and deserialized as an attribute, and not as a nested element. For example:

**XML:**

```
<xsd:complexType name="Sample2">
  <xsd:sequence>
    <xsd:attribute name="a" type="xsd:string"/>
  </xsd:sequence>
</xsd:complexType>
```

**Java:**

```

public class Sample2 {
    // ..
    public Sample2() {}

    // Bean Property a
    public String getA() {...}
    public void setA(String value) {...}
}

```

– *Qualified versus unqualified names*

The **WSDL2Java** command supports the `elementForm` and `attributeForm` schema attributes.

This support is not specified in the JAX-RPC specification. These attributes are used to indicate whether an element or attribute is serialized and deserialized with a qualified or unqualified name. The default setting for `elementForm` is `qualified` and the default setting for `attributeForm` is `unqualified`. These settings do not affect the Java bean class that is generated, but the information is captured in the `_Helper` class metadata.

– *Extension and the abstract attribute*

The **WSDL2Java** command supports extension of an `xsd:complexType` through the `xsd:extension` element. This support is required by the JAX-RPC specification.

The **WSDL2Java** command supports the abstract attribute. This feature is listed as optional by the JAX-RPC specification.

The following example shows the accepted use of the extension and abstract constructs. WebSphere Application Server uses the extension and abstract constructs to support polymorphism.

**XML:**

```

<xsd:complexType name="Base" abstract="true">
  <xsd:sequence>
    <xsd:element name="a" type="xsd:int" />
  </xsd:sequence>
</xsd:complexType>

<xsd:complexType name="Derived">
  <xsd:complexContent>
    <xsd:extension base="ns:Base">
      <xsd:sequence>
        <xsd:element name="b" type="xsd:int" />
      </xsd:sequence>
    </xsd:extension>
  </xsd:complexContent>
</xsd:complexType>

```

**Java:**

```

public abstract class Base {
    // ...
    public Base() {}

    public int getA() {...}
    public void setA(int a) {...}
}

public class Derived extends Base {
    // ...
    public Derived() {}
}

```

```

    public int getB() {...}
    public void setB(int b) {...}
}

```

– *Support for xsd:any*

The **WSDL2Java** command supports `xsd:anyElement`, which is different than `xsd:anyType`. This feature is not defined within the JAX-RPC specification and is subject to change.

If an `<xsd:any/>` element is defined within `xsd:sequence` or `xsd:all` group, SOAP values that do match one of the `xsd:elements` are stored in the Java bean as `com.ibm.ws.webservices.engine.xmlsoap.SOAPElement` objects. Values can be accessed from the Java bean using the `get_any()` and `set_any()` methods.

• **Mapping of xsd:element**

An `xsd:element` is a construct that has a name or name attribute, and a type defined by a `complexType` or primitive type. There are two different kinds of `xsd:elements`:

- **Root:** Defined directly underneath the schema elements and referenced by other constructs.
- **Nested:** Nested underneath group elements and are not referenced by other constructs.

Root elements are referenced by the WSDL file constructs, especially if the WSDL file is used to describe a literal service. Typically, root elements and types have the same names, which is allowed in the schema language. Under most circumstances the **WSDL2Java** command can produce Java artifacts without name collisions.

– *Four ways to reference a type*

There are four ways that a nested or root `xsd:element` can reference a type:

- **Use the type attribute:**

This is the most common way to reference a type, for example:

```
<xsd:element name="one" type="ns:myType" />
```

The **WSDL2Java** command recognizes the type attribute as a reference to a `complexType` or `simpleType` named, `myType`. The **WSDL2Java** command generates a Java type based on the characteristics of `myType`. Support for the type attribute is required by the JAX-RPC specification.

- **Use the ref attribute:** For example:

```
<xsd:element ref="ns:myElement" />
```

The **WSDL2Java** command recognizes the `ref` attribute as a reference to another root element named `myElement`. The name of the element is obtained from the referenced element, such as `myElement`. The type of the element is the type of the referenced element. The **WSDL2Java** command generates a Java type based on the characteristics of the referenced type. The `ref` attribute is an optional feature of the JAX-RPC specification.

- **Use no attribute:**

For example:

```
<xsd:element name="three" />
```

When you do not use an attribute, the **WSDL2Java** command recognizes a reference to the `xsd:anyType` as defined by the XML schema specification. The `xsd:anyType` is an optional type of the JAX-RPC specification.

- **Use an anonymous type:**

For example:

```
<xsd:element name="four">
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="foo" type="xsd:string" />
    </xsd:sequence>
  </xsd:complexType>
</xsd:element>
```

When you use an anonymous type, the **WSDL2Java** command recognizes a reference to the type defined within the element.

**Note:** The `complexType` does not have a name. The **WSDL2Java** command generates a Java type based on the characteristics of this type. Since the anonymous type does not have a name, the **WSDL2Java** command uses the name of the container element, which can result in collisions with defined types and other anonymous types. The **WSDL2Java** command automatically detects and renames classes to avoid collisions. Support for anonymous types is not defined by the JAX-RPC specification, however using anonymous types is common.

**Note:** An `xsd:attribute` is like an `xsd:element`; it contains a name and refers to a type. An `xsd:attribute` can refer to its type with the type attribute or using an anonymous type.

- *Element specific attributes*

Some attributes can be applied to `xsd:elements` and not to XML types.

The `maxOccurs` attribute indicates the maximum number of occurrences of the element in the SOAP message. The default value is 1. If the value is greater than 1, or unbounded, the **WSDL2Java** command maps the construct to a Java array or bean indexed property. Metadata is also generated to properly serialize and deserialize a series of elements versus a normal XML array. The `maxOccurs` attribute is an optional feature of the JAX-RPC specification.

The `minOccurs` attribute indicates the minimum number of occurrences of the element in the SOAP message. The default value is 1. The `xsi:nilable` attribute indicates whether the element can have a `nil` value. The `minOccurs` and `xsi:nilable` settings affect how a `null` value is serialized in a SOAP message. If `minOccurs=0`, the `null` value is not serialized. If `xsi:nilable=true`, the value is serialized with the `xsi:nil=true` attribute.

• **Mapping of `xsd:complexType` to Java array**

The **WSDL2Java** command maps the following three kinds of XML formats to Java arrays:

**XML:**

```
<xsd:element name="array1" type="soapenc:Array" />
```

**Java:**

```
Object[] array1;
```

**XML:**

```
<xsd:complexType name="arrayOfInt">
  <xsd:complexContent>
```

```

<xsd:restriction base="soapenc:Array">
  <xsd:attribute ref="soapenc:arrayType" wsdl:arrayType="xsd:int[]" />
</xsd:restriction>
</xsd:complexContent>
</xsd:complexType>

```

```

<xsd:element name="array2" type="ns:arrayOfInt" />

```

**Java:**

```

int[] array2;

```

**XML:**

```

<xsd:complexType name="arrayOfInt">
  <xsd:complexContent>
    <xsd:restriction base="soapenc:Array">
      <xsd:sequence>
        <xsd:element name="item" type="xsd:int" maxOccurs="unbounded" />
      </xsd:sequence>
    </xsd:restriction>
  </xsd:complexContent>
</xsd:complexType>

```

```

<xsd:element name="array3" type="ns:arrayOfInt" />

```

**Java:**

```

int[] array3;

```

- **Mapping of xsd:simpleType enumeration**

The **WSDL2Java** command maps the following XML enumeration to a JAX-RPC specified enumeration class. See section 4.2.4 of the JAX-RPC specification for more details.

```

<xsd:simpleType name="EyeColorType" >

  <xsd:restriction base="xsd:string">
    <xsd:enumeration value="brown"/>
    <xsd:enumeration value="green"/>
    <xsd:enumeration value="blue"/>
  </xsd:restriction>
</xsd:simpleType>

```

- **Mapping of xsd:complexType to exception class**

If a complexType is referenced in a wsdl:message for a wsdl:fault, the complexType is mapped to a class that extends the exception, java.lang.Exception. This mapping is similar to the mapping of a complexType to a Java bean class, except a full constructor is generated, and only getter methods are generated. See section 4.3.6 of the JAX-RPC specification for more details.

- **Other mappings**

The **WSDL2Java** command supports the mapping of xsd:simpleType and xsd:complexTypes that extend xsd:simpleTypes. These constructs are mapped to Java bean classes. The simple value is mapped to a Java bean property named, value. This is an optional feature of the JAX-RPC specification.

### Mapping of wsdl:portType

The wsdl:portType construct is mapped to the Service Endpoint Interface. The name of the wsdl:portType is mapped to the name of the class of the Service Endpoint Interface.

### Mapping of wsdl:operation

A `wsdl:operation` within a `wsdl:portType` is mapped to a method of the Service Endpoint Interface. The name of the `wsdl:operation` is mapped to the name of the method. The `wsdl:operation` contains `wsdl:input` and `wsdl:output` elements that reference the request and response `wsdl:message` constructs using the `message` attribute. The `wsdl:operation` can contain a `wsdl:fault` element that references a `wsdl:message` describing the fault. These faults are mapped to Java classes that extend the exception, `java.lang.Exception` as discussed in section 4.3.6 of the JAX-RPC specification.

- **Effect of document literal wrapped format**

If the WSDL file uses the .NET document and literal wrapped format, the method parameters are mapped from the wrapper `xsd:element`. The .NET document and literal format is automatically detected by the **WSDL2Java** command. The following criteria must be met:

- The WSDL file must have `style="document"` in its `wsdl:binding` constructs.
- The WSDL file must have `use="literal"` in its `wsdl:binding` constructs.
- The `wsdl:message` referenced by the `wsdl:operation` input construct must have a single part.
- The part must use the `element` attribute to reference an `xsd:element`.
- The referenced `xsd:element`, or wrapper element, must have the same name as the `wsdl:operation`.
- The wrapper element must not contain any `xsd:attributes`.

In such cases, each parameter name is mapped from a nested `xsd:element` contained within wrapper element. The type of the parameter is mapped from the type of the nested `xsd:element`. For example:

**XML:**

```
<xsd:element name="myMethod" >
  <xsd:complexType>
    <xsd:sequence>
      <xsd:element name="param1" type="xsd:string" />
      <xsd:element name="param2" type="xsd:int" />
    </xsd:sequence>
  </xsd:complexType>
</xsd:element>
...
<wsdl:message name="response" />
  <part name="parameters" element="ns:myMethod" />
</wsdl:message name="response" />

<wsdl:message name="response" />
...
<wsdl:operation name="myMethod">
  <input name="input" message="request" />
  <output name="output" message="response" />
</wsdl:operation>
```

**Java:**

```
void myMethod(String param1, int param2) ...
```

- **Parameter mapping**

If the document and literal wrapped format is not detected, the parameter mapping follows the normal JAX-RPC mapping rules set in section 4.3.4 of the JAX-RPC specification.

Each parameter is defined by a `wsdl:message` part referenced from the input and output elements.

- A `wsdl:part` in the request `wsdl:message` is mapped to an input parameter.

- A `wsdl:part` in the response `wsdl:message` is mapped to the return value. If there are multiple `wsdl:parts` in the response message, they are mapped to output parameters.
  - A `Holder` class is generated for each output parameter as discussed in section 4.3.5 of the JAX-RPC specification.
- A `wsdl:part` that is both the request and response `wsdl:message` is mapped to an `inout` parameter.
  - A `Holder` class is generated for each `inout` parameter as discussed in section 4.3.5 of the JAX-RPC specification.
  - The `wsdl:operation parameterOrder` attribute defines the order of the parameters.

The **WSDL2Java** command supports overloaded methods, but confirm that the part names of the overloaded methods are unique. For example:

**XML:**

```
<wsdl:message name="request" >
  <part name="param1" type="xsd:string" />
  <part name="param2" type="xsd:int" />
</wsdl:message name="response" />

<wsdl:message name="response" />
...
<wsdl:operation name="myMethod" parameterOrder="param1, param2">
  <input name="input" message="request" />
  <output name="output" message="response" />
</wsdl:operation>
```

**Java:**

```
void myMethod(String param1, int param2) ...
```

**Mapping of `wsdl:binding`**

The **WSDL2Java** command uses the `wsdl:binding` information to generate an implementation specific client side stub. WebSphere Application Server uses the `wsdl:binding` information on the server side to properly deserialize the request, invoke the Web service, and serialize the response. The information in the `wsdl:binding` should not affect the generation of the Service Endpoint Interface, but it can when the document and literal wrapped format is used or when there are MIME attachments.

• **MIME attachments**

For a WSDL 1.1 compliant WSDL file, a part of an operation message, which is defined in the binding to be a MIME attachment, becomes a parameter of the type of the attachment regardless of the part declared. For example:

**XML:**

```
<wsdl:types>
  <schema ...>
    <complexType name="ArrayOfBinary">
      <restriction base="soapenc:Array">
        <attribute ref="soapenc:arrayType" wsdl:arrayType="xsd:binary[]" />
      </restriction>
    </complexType>
  </schema>
</wsdl:types>

<wsdl:message name="request">
  <part name="param1" type="ns:ArrayOfBinary" />
<wsdl:message name="response" />
```



```

<wsdl:message name="response" />
...

<wsdl:operation name="myMethod">
  <input name="input" message="request" />
  <output name="output" message="response" />
</wsdl:operation>
...

<binding ...
  <wsdl:operation name="myMethod">
    <input>
      <mime:multipartRelated>
        <mime:part>
          <mime:content part="param1" type="image/jpeg"/>
        </mime:part>
      </mime:multipartRelated>
    </input>
    ...
  </wsdl:operation>

```

**Java:**

```
void myMethod(java.awt.Image param1) ...
```

The JAX-RPC requires support for the following MIME types:

| MIME type       | Java type                         |
|-----------------|-----------------------------------|
| image/gif       | java.awt.Image                    |
| image/jpeg      | java.awt.Image                    |
| text/plain      | java.lang.String                  |
| multipart/*     | javax.mail.internet.MimeMultipart |
| text/xml        | javax.xml.transform.Source        |
| application/xml | javax.xml.transform.Source        |

There are a number of problems with MIME attachments as they are defined in WSDL 1.1, including:

- The semantics of the mime:multipartRelated clause are not fully defined
- The semantics do not allow for arrays of MIME attachments

Because of these problems, several types are not specified by the JAX-RPC for MIME attachments. These types are defined in the supported mappings previously discussed.

- **Headers**

A wsdl:binding can also define SOAP headers, for example:

**XML:**

```

<wsdl:message name="request">
  <part name="param1" type="xsd:string" />
</wsdl:message/>

<wsdl:message name="response" />
...

<wsdl:operation name="myMethod">
  <input name="input" message="request" />
  <output name="output" message="response" />
</wsdl:operation>
...

```

```

<binding ...
  <wsdl:operation name="myMethod">
    <input>
      <soap:header message="request" part="param1" use="literal" />
    </input>
    ...
  </wsdl:operation>

```

**Java:**

```
void myMethod(String param1) ...
```

This is an example of an explicit header or a header with a value determined from a method parameter. Instead of appearing in the soap:body SOAP message, the value of param1 now appears in the soap:header SOAP message. The **WSDL2Java** command supports explicit headers and does not support implicit headers. Implicit headers have a value not determined by a parameter. For example, you could replace the soap:header clause in the example with:

```
<soap:header message="someOtherMsgNotAppearingInThePortType" part="someOtherPart" use="literal"/>
```

**Note:** The **WSDL2Java** command supports explicit headers, but it is not considered good programming practice to use them. Headers are typically used for middleware logic, not business logic. Explicit headers place parameters used in business logic into the header.

**Mapping of wsdl:service**

The wsdl:service element is mapped to a Generated Service interface. The Generated Service interface contains methods to access each of the ports in the wsdl:service. The Generated Service interface is discussed in sections 4.3.9, 4.3.10, and 4.3.11 of the JAX-RPC specification.

In addition, the wsdl:service element is mapped to the implementation-specific ServiceLocator class, which is an implementation of the Generated Service interface.

**Mapping between WSDL and SOAP messages**

The WSDL file defines the format of the SOAP message that is sent over the wire. The **WSDL2Java** command and the WebSphere Application Server run time use the information in the WSDL file to confirm that the SOAP message is properly serialized and deserialized.

**Document versus RPC, literal versus encoded**

If a wsdl:binding indicates a message is sent using an RPC format, the SOAP message contains an element defining the operation. If a wsdl:binding indicates the message is sent using a document format, the SOAP message does not contain the operation element.

If the wsdl:part is defined using the type attribute, the name and type of the part are used in the message. If the wsdl:part is defined using the element attribute, the name and type of the element are used in the message. The element attribute is not allowed by the JAX-RPC specification when use="encoded".

If a `wSDL:binding` indicates a message is encoded, the values in the message are sent with `xsi:type` information. If a `wSDL:binding` indicates that a message is literal, the values in the message are typically not sent with `xsi:type` information. For example:

**WSDL:**

```
<xsd:element name="c" type="xsd:int" />
...
<wsdl:message name="request">
  <part name="a" type="xsd:string" />
  <part name="b" element="ns:c" />
</wsdl:message>
...
<wsdl:operation name="method" >
  <input message="request" />
...

```

**RPC/ENCODED:**

```
<soap:body>
  <ns:method>
    <a xsi:type="xsd:string">ABC</a>
    <element attribute is not allowed in rpc/encoded mode>
  </ns:method>
</soap:body>

```

**DOCUMENT/LITERAL:**

```
<soap:body>
  <a>ABC</a>
  <c>123</a>
</soap:body>

```

**DOCUMENT/LITERAL wrapped:**

```
<soap:body>
  <ns:method_wrapper>
    <a>ABC</a>
    <c>123</a>
  </ns:method_wrapper>
</soap:body>

```

The document and literal wrapped mode is the same as the document and literal mode. However, in the document and literal wrapped mode, there is only a single element within the body, and the element has the same name as the operation.

### Multi-ref processing

If `use=encoded`, XML types that are not simpleTypes are passed in the SOAP message using the multi-ref attributes, `href` and `id`. The following example assumes that parameters one and two reference the same Java bean named, `info` containing fields `a` and `b`:

**Note:**

Deserialization produces a single instance of the `info` class for the encoded case and two instances are created for the literal case.

**RPC/ENCODED:**

```
<soap:body>
  <ns:method>
    <param1 href="#id1" />
    <param2 href="#id2" />
  </ns:method>
  <multiref id="id1" xsi:type="ns:info">

```

```

    <a xsi:type="xsi:string">hello</a>
    <b xsi:type="xsi:string">world</b>
  </multiref>
</soap:body>

```

**RPC/LITERAL:**

```

<soap:body>
  <ns:method>
    <param1>
      <a>hello</a>
      <b>world</b>
    </param1>
    <param2>
      <a>hello</a>
      <b>world</b>
    </param2>
  </ns:method>
</soap:body>

```

### XML arrays and the maxOccurs attribute

A SOAP message is affected by whether the element is defined by an XML array or using the maxOccurs attribute.

**WSDL:**

```

<element name="foo" type="ns:ArrayOfString" />

```

**Literal Instance:**

```

<foo>
  <item>A</item>
  <item>B</item>
  <item>C</item>
</foo>

```

**WSDL:**

```

<element name="foo" maxOccurs="unbounded" type="xsd:string"/>

```

**Literal Instance:**

```

<foo>A</foo>
<foo>B</foo>
<foo>C</foo>

```

### minOccurs and nillable attributes

An element specified with minOccurs=0 that has a null value is not serialized in the SOAP message. An element specifying nillable="true" has a null value and is serialized into a SOAP message with the xsi:nil=true attribute. For example:

```

<a xsi:nil="true" />

```

### Qualified versus unqualified

The XML schema attributeForm and elementForm attributes indicate whether the attributes and nested elements are serialized with qualified or unqualified names. If a part name is serialized, it is always serialized as an unqualified name.

## Developing a Web services client

### Before you begin

Locate the Web Services Description Language (WSDL) file that defines the Web service to access.

To create the client code and artifacts that enable the application client to access a Web service:

#### Steps for this task

1. Develop implementation templates and bindings from a WSDL file.  
Follow the steps in this task, but use the `-role client -container client` option with the **WSDL2Java** command tool. The client-side bindings and artifacts are generated.
2. Complete the client implementation.
3. (Optional) Configure the `webservicesclient.xml` deployment descriptor.  
Complete this step if you are developing a managed client that runs in the J2EE client container.
4. (Optional) Configure the `ibm-webservices-bnd.xmi` deployment descriptor.  
Complete this step if you are deploying a managed client that runs in the J2EE client container and you want to override the default client settings. See Web services assembly properties for more information about the `ibm-webservicesclient-bnd.xmi` deployment descriptor.
5. (Optional) Assemble a Web services-enabled client Java archive (JAR) file.  
Complete this step if you are developing a managed client that runs in the J2EE client container.
6. Test the Web services-enabled client application.

#### Results

You have created and tested a Web services client application.

## **Assembling a Web services-enabled client JAR and EAR file**

#### Before you begin

You need the following artifacts:

- Assembled client module, containing the implementation, all classes generated by the **WSDL2Java** command, `MANIFEST.MF` and deployment descriptor. This module can be:
  - An application client module containing `META-INF/application-client.xml`
  - A Web module containing `WEB-INF/web.xml`
  - An enterprise JavaBean (EJB) module containing `META-INF/ejb-jar.xml`
- Web Services Description Language (WSDL) file used to develop the client
- Configured `webservicesclient.xml` and `ibm-webservicesclient-bnd.xmi`(if used) deployment descriptors
- Generated JAX-RPC mapping deployment descriptor

The steps in this topic explain how to use the command-line tools to assemble a Web service-enabled client application. You can also use the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>) to assemble Web service-enabled client applications.

To create the client code and artifacts that enable the application client to access a Web service:

#### Steps for this task

1. For a Web services client in an application client Java archive (JAR) or EJB JAR file, arrange the following files:
  - a. Expand the JAR file in the *root* directory.
  - b. Place the WSDL file in the META-INF/wsdl subdirectory.  
The WSDL file is indicated by the <wsdl-file> element in the webservicescient.xml file.
  - c. Place the webservicescient.xml and the JAX-RPC mapping file in the META-INF subdirectory.  
The JAX-RPC mapping file is indicated by the <jaxrpc-mapping-file> element in the webservicescient.xml file.
  - d. (Optional) Place the ibm-webservicescient-bnd.xmi file in the META-INF subdirectory, if used.
2. Add these files to the existing JAR file with the appropriate command for your platform.  
For UNIX platforms use **jar -uvf existing.jar META-INF/\***. For Windows platforms use **jar -uvf existing.jar META-INF\\***
3. Assemble the JAR file into an EAR file using typical assembly techniques if the client runs in a container.
4. For a Web services client in a Web archive (WAR) file, arrange the following files:
  - a. Expand the WAR file in the *root* directory.
  - b. Place the WSDL file in the WEB-INF/wsdl subdirectory.
  - c. Place the webservicescient.xml and JAX-RPC mapping files in the WEB-INF subdirectory.
  - d. (Optional) Place the ibm-webservicescient-bnd.xmi file in the WEB-INF subdirectory, if used.
5. Add these files to the existing WAR file with the appropriate command for your platform.  
For UNIX platforms use **jar -uvf existing.war WEB-INF/\***. For Windows platforms use **jar -uvf existing.war WEB-INF\\***

### Results

The artifacts required to enable the client module to use Web services for J2EE are added to the JAR file.

### Usage scenario

This example uses a JAR file named AddressBookClient.jar and an EAR file named AddressBookClient.ear:

After running the **jar -u** command, the AddressBookClient.jar file contains the following files. The files added in this task are in bold:

```
META-INF/MANIFEST.MF
META-INF/application-client.xml
META-INF/wsdl/AddressBook.wsdl
META-INF/webservicescient.xml
META-INF/AddressBook_mapping.xml
com/ibm/websphere/samples/webservices/addr/Address.class
com/ibm/websphere/samples/webservices/addr/AddressBook.class
com/ibm/websphere/samples/webservices/addr/AddressBookClient.class
com/ibm/websphere/samples/webservices/addr/AddressBookService.class
...other generated classes...
```

After assembling the `AddressBookClient.jar` file into the `AddressBookClient.ear` file, the `AddressBookClient.ear` file contains the following files:

`META-INF/MANIFEST.MF`  
`AddressBookClient.jar`  
`META-INF/application.xml`

### What to do next

Test the client JAR or EAR file.

## Testing Web services-enabled clients

### Before you begin

Before testing your Java client, confirm that the server endpoint specified in the client Web Services Description Language (WSDL) file is running and available.

The following steps and examples assume that you are testing a system that has WebSphere Application Server installed, and that the `WAS_HOME` environment variable is set in the WebSphere Application Server installation directory.

Tests are run differently depending on whether the client module has client container deployment information, which consists of the `application-client.xml` and `webservicesclient.xml` files, as well as the JAX-RPC mapping file and WSDL file. The client enterprise archive (EAR) files discussed in this topic are referred to as managed because they contain the deployment information. The client Java archive (JAR) files discussed are referred to as unmanaged because they do not contain the deployment information.

To test Web services-enabled clients:

### Steps for this task

1. Test an unmanaged client JAR file.
  - a. Set your `CLASSPATH` when you Set up a Web services development environment.
  - b. Add your client JAR file to `CLASSPATH`.
  - c. Execute your application with the `java` command.

The unmanaged client application runs.

2. Test a managed client EAR file.
  - a. Execute your client application with the `launchClient` command.

An example of using the command is as follows:

```
launchClient clientEar
```

### Results

Web services-enabled clients that have been tested.

### What to do next

Troubleshoot your Web services application.

---

## Assembling Web services applications

### Before you begin

You must assemble either an enterprise JavaBean (EJB) Java archive (JAR) file or an EJB Web archive (WAR) file before assembling the Web services-enabled enterprise archive (EAR) file.

The steps in this topic explain how to use the command-line tools to assemble Web service J2EE modules. You can also assemble Web service J2EE modules with the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>).

To assemble Web services applications:

#### Steps for this task

1. Assemble a Web services-enabled EJB JAR file.
2. Assemble a Web services-enabled WAR file.
3. Access a Web Services Description Language (WSDL) file
4. Assemble a Web services-enabled EAR file.
5. (Optional) Enable the EAR file.

If the EAR file contains EJB modules, it must have the Web services endpoint WAR file added with the **endptEnabler** command before deployment.

#### Results

A Web services-enabled EAR file that you can deploy into WebSphere Application Server.

#### What to do next

Deploy the Web services-enabled EAR file into WebSphere Application Server.

## **Assembling a Web services-enabled EJB JAR file**

You can assemble a Web services-enabled enterprise JavaBean (EJB) Java archive (JAR) file in one of two ways:

#### Steps for this task

1. Assemble a Web services-enabled EJB JAR file when starting from Java code
2. Assemble a Web services-enabled EJB JAR file when starting from Web Services Description Language (WSDL)

#### Results

An assembled Web services-enabled EJB JAR file.

#### What to do next

Deploy Web services.

## **Assembling a Web services-enabled EJB JAR file when starting from Java code**

### Before you begin

You need the following artifacts:

- Assembled enterprise JavaBean (EJB) Java archive (JAR) file (not enabled for Web services)



- Compiled Java class for the Service Endpoint Interface
- Web Services Description Language (WSDL) file
- Compiled Java classes for all classes emitted by the **WSDL2Java** command tool. A compiled Java class is a result of developing an enterprise bean starting with WSDL.
- Complete `webservices.xml`, `ibm-webservices-bnd.xmi` and Java API for XML-based remote procedure call (JAX-RPC) mapping deployment descriptors.

The steps in this topic explain how to use the command-line tools to assemble a Web service-enabled EJB JAR file. You can also use the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>) to assemble a Web service-enabled EJB JAR file.

To assemble an Web services-enabled EJB JAR file when starting from Java code:

#### Steps for this task

1. Place the JAR file in the *root* directory.
2. Place the WSDL file specified by the deployment descriptor `<wsdl-file>` element in the `META-INF/wsdl` subdirectory.
3. Place the JAX-RPC mapping file as specified by the deployment descriptor `<jaxrpc-mapping-file>` element in the `META-INF` subdirectory.
4. Place the `webservices.xml` and `ibm-webservices-bnd.xmi` files in the `META-INF` subdirectory.
5. Place the Service Endpoint Interface class in a subdirectory corresponding to its Java package.
6. Add these files to the EJB JAR file.

Assuming that the Service Endpoint Interface class is under the `com` subdirectory, and the WSDL file is in the `META-INF/wsdl` subdirectory, execute the **`jar -uvf JAR_file com META-INF/*`** command.

#### Results

The artifacts required to Web service-enable an EJB module for Web services are added to the JAR file.

#### Usage scenario

After running the **`jar -u`** command on a JAR file named `AddressBook.jar`, the JAR file contains the following files. The files added in this task are in bold:

```

META-INF/MANIFEST.MF
META-INF/ejb-jar.xml
addr/Address.class
addr/AddressBook_RI.class
addr/AddressBookBean.class
addr/AddressBookHome.class
addr/Phone.class
addr/StateType.class
addr/AddressBook.class
META-INF/wsdl/AddressBook.wsdl
META-INF/ibm-webservices-bnd.xmi
META-INF/webservices.xml
META-INF/AddressBook_mapping.xml

```

#### What to do next

Assemble a Web services-enabled EAR file.

## Assembling Web services-enabled EJB JAR file when starting from Web Services Description Language

Before you begin

You need the following artifacts:

- An assembled enterprise JavaBean (EJB) Java archive (JAR) file containing the EJB implementation and all classes generated by the **WSDL2Java** command tool when the **role** argument is `develop-server` and the **container** argument is `EJB`.
- A Web Services Description Language (WSDL) file
- Complete `webservices.xml`, `ibm-webservices-bnd.xmi`, and Java API for XML-based remote procedure call (JAX-RPC) mapping deployment descriptors.

This topic explains how to assemble a Web services-enabled JAR file with command-line tools. You can also use the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>) to assemble Web services-enabled JAR files.

To assemble a Web services-enabled EJB JAR file when starting from WSDL:

### Steps for this task

1. Place the JAR file in the `root <samp/>` directory.
2. Place the WSDL file as specified by the deployment descriptor `<samp><wsdl-file>` element in the `META-INF/wsdl` subdirectory.
3. Place the JAX-RPC mapping file as specified by the deployment descriptor `<jaxrpc-mapping-file>` element in the `META-INF` subdirectory.
4. Place the `webservices.xml` and `ibm-webservices-bnd.xmi` files in the `META-INF` subdirectory.
5. Add these files to the EJB JAR file.

Assuming that the WSDL file is in the `META-INF` subdirectory, execute the `jar -uvf EAR_file META-INF/*` command.

### Results

The artifacts required to enable an EJB module for Web services are added to the JAR file.

### Usage scenario

After running the `jar -u` command, a JAR file named `AddressBook.jar` contains the following files. The files added in this task are in bold:

```
META-INF/MANIFEST.MF
META-INF/ejb-jar.xml
addr/Address.class
addr/AddressBook_RI.class
addr/AddressBookSoapBindingImpl.class
addr/AddressBookHome.class
addr/Phone.class
addr/StateType.class
addr/AddressBook.class
META-INF/wsdl/AddressBook.wsdl
META-INF/ibm-webservices-bnd.xmi
META-INF/webservices.xml
META-INF/AddressBook_mapping.xml
```

### What to do next

Assemble a Web services-enabled enterprise archive (EAR) file.

## **Assembling Web services-enabled WAR file**

You can assemble a Web services-enabled Web archive (WAR) file in one of two ways:

### Steps for this task

1. Assemble a Web services-enabled WAR file when starting from Java code.
2. Assemble a Web services-enabled WAR file when starting from Web Services Description Language (WSDL).

### Results

A Web services-enabled WAR file is assembled.

### What to do next

Assemble a Web services-enabled EAR file

## **Assembling a Web services-enabled WAR file when starting from Java code**

### Before you begin

You need the following artifacts:

- An assembled Web archive (WAR) file containing `web.xml`, but not Web services-enabled
- Compiled Java class for the Service Endpoint Interface
- A Web Services Description Language (WSDL) file
- Complete `webservices.xml`, `ibm-webservices-bnd.xmi` and Java API for XML-based remote procedure call (JAX-RPC) mapping deployment descriptors.

The steps in this topic explain how to use the command-line tools to assemble a Web service-enabled WAR file. You can also use the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>) to assemble a Web service-enabled WAR file.

To assemble a Web services-enabled WAR file when starting from Java code:

### Steps for this task

1. Expand the WAR file into a directory.
2. Confirm that the `WEB-INF/web.xml` descriptor for the Web module contains a `<servlet-class>` element indicating the Java bean class that is implementing the service.
3. Place the WSDL file as specified by the deployment descriptor `<wsdl-file>` element of `webservices.xml` file in the `WEB-INF/wsdl` subdirectory.
4. Place the JAX-RPC mapping file as specified by the deployment descriptor `<jaxrpc-mapping-file>` element of `webservices.xml` in the `WEB-INF` subdirectory.
5. Place the `webservices.xml` and `ibm-webservices-bnd.xmi` deployment descriptors in the `WEB-INF` subdirectory.

6. Place the Service Endpoint Interface class in a subdirectory corresponding to its Java package.
7. Add these files to the WAR file.  
Assuming that the Service Endpoint Interface class is under the `com` subdirectory, and the WSDL and mapping files are in `WEB-INF`, execute the `jar -uvf WAR_file com WEB-INF/*` command.

### Results

The artifacts required to Web service-enable the Web module are added to the WAR file.

### What to do next

Assemble a Web services-enabled enterprise archive (EAR) file

## **Assembling a Web services-enabled WAR file when starting from Web Services Description Language**

### Before you begin

You need the following artifacts:

- Assembled Web archive (WAR) file containing the enterprise JavaBean (EJB) implementation, all classes generated by the **WSDL2Java** command tool and a Web deployment descriptor, `web.xml`
- A Web Services Description Language (WSDL) file
- Complete `webservices.xml`, `ibm-webservices-bnd.xmi`, and Java API for XML-based remote procedure call (JAX-RPC) mapping deployment descriptors.

This topic explains how to assemble a Web services-enabled WAR file with command-line tools. You can also use the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>) to assemble Web services-enabled WAR files.

To assemble a Web services-enabled WAR file when starting from WSDL:

### Steps for this task

1. Expand the WAR file into a directory.
2. Confirm that the `WEB-INF/web.xml` deployment descriptor for the Web module contains a `<servlet>` element including the `<servlet-name>` element.  
The `<servlet-name>` element can be any string and `<servlet-class>` element indicating the Java bean class that is implementing the service.
3. Place the WSDL file as specified by the `webservices.xml` deployment descriptor `<wsdl-file>` element in the `WEB-INF/wsdl` subdirectory.
4. Place the JAX-RPC mapping file as specified by the `webservices.xml` deployment descriptor `<jaxrpc-mapping-file>` element in the `WEB-INF` subdirectory.
5. Place the `webservices.xml` and `ibm-webservices-bnd.xmi` deployment descriptors in the `WEB-INF` subdirectory.
6. Add these files to the WAR file.  
Assuming that the WSDL and mapping files are in `WEB-INF`, execute the `jar -uvf WAR_file com WEB-INF/*` command.

### Results

The artifacts required to Web service-enable the Web module is added to the WAR file.

### What to do next

Assemble a Web services-enabled enterprise archive (EAR) file

## **Assembling a Web services-enabled EAR file**

### Before you begin

Before assembling a Web services-enabled enterprise archive (EAR) file, one of these three tasks should be complete:

- Assemble a Web services-enabled EJB Java archive (JAR) file.
- Assemble a Web services-enabled EJB Web archive (WAR) file.

This topic explains how to assemble a Web services-enabled EAR file with command-line tools. You can also assemble a Web services-enabled EAR file with the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>).

To assemble a Web services-enabled EAR file:

### Steps for this task

1. Assemble the Web services-enabled JAR or WAR file into an EAR file.

Now assemble the EAR file that contains the JAR or WAR files. The EAR file can contain an enterprise bean or application client JAR files; Web applications or WAR files; and metadata describing the applications or `application.xml` files.

### Results

A Web services-enabled EAR file.

### Usage scenario

In the following example, there is an `application.xml` deployment descriptor packaged with a Web services-enabled JAR file called `AddressBook.jar` that is packaged into an EAR file called `AddressBook.ear`. The EAR file contains:

```
META-INF/MANIFEST.MF
META-INF/application.xml
AddressBook.jar
```

An example of the `application.xml` deployment descriptor is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<!DOCTYPE application PUBLIC "-//Sun Microsystems, Inc.//DTD J2EE Application 1.3//EN"
    "http://java.sun.com/dtd/application_1_3.dtd">
<application id="Application_ID">
  <display-name>AddressBookJ2WEE</display-name>
  <description>AddressBook EJB Example from Java</description>
  <module id="EjbModule_1">
    <ejb>AddressBook.jar</ejb>
  </module>
</application>
```

### What to do next

If the EAR file contains EJB modules, Run the endptEnabler. Then, deploy the EAR file into WebSphere Application Server.

## Web services assembly properties

### ibm-webservices-bnd.xmi properties

The `ibm-webservices-bnd.xmi` file is a deployment descriptor for a Web Services-enabled Web module or enterprise JavaBean (EJB) module. It contains information for the Web services runtime that is either WebSphere product-specific or was not specified by Web services for J2EE.

You can edit these properties using the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>):

1. Locate the `webservices.xml` file in the module.
2. Double-click the `webservices.xml` file to open the **Web Services** editor.
3. Access the **Web Services Bindings** editor through the **Bindings** tab at the bottom of the editor pane.
4. Access the **Web Services Binding Configurations** editor through the **Binding Configurations** tab at the bottom of the editor pane.
5. After editing the properties, type **ctrl-s** on your keyboard to save the changes.

The following user-definable assembly properties are supported:

- **wsDescNameLink**

Attribute of the `wsdescBindings` element that specifies the link to the corresponding `<webservice-description-name>` in `webservices.xml`.

You can use the Assembly Toolkit to set this property:

1. Open the **Web Service Bindings** editor.
2. Expand the **Web Service Description Bindings** section.
3. Click **Add** and choose the Web services description binding properties for which you want to apply the change.
4. Click **OK**.

- **pc-name-link**

Attribute of the `pcBindings` element that specifies the link to the `<port-component-name>` in the `webservices.xml` file.

You can edit these properties in the Assembly Toolkit:

1. Open the **Web Services Bindings** editor.
2. Expand the **Port Component Binding** section.
3. Click **Add**.
4. Select the `port_component_name` from the drop down list in the **PC Name Link** field.

- **scope**

Attribute of the `pcBindings` element that specifies when new instances of implementation beans are created. Possible values are Request, Session, and Application.

You can edit these properties in the Assembly Toolkit:

1. Open the **Web Services Bindings** editor.
2. Expand the **Port Component Binding** section.
3. Click **Add**.

4. Select the *implementation\_scope\_name* from the drop down list in the **Scope** field.

The value of scope for a deployed Web service can be changed using the administrative console. Using application management, navigate to the Web module of the Web service application and select Web Services Implementation Scope.

### **ibm-webservicesclient-bnd.xml properties**

The `ibm-webservicesclient-bnd.xml` file contains information for the Web Services runtime that is WebSphere product-specific.

You can edit these properties using the Assembly Toolkit:

1. Locate the `webservicesclient.xml` file in the module.
2. Double-click the `webservices.xml` file to open the **Web Services Client** editor.
3. Access the **Web Services Client Bindings** editor through the **Client Binding** tab at the bottom of the editor pane.
4. Access the **Web Services Client Port Bindings** editor through the **Port Bindings** tab at the bottom of the editor pane.
5. After editing the properties, type **ctrl-s** on your keyboard to save the changes.

### **Assembly properties**

The following user-definable assembly properties are supported:

- **componentNameLink**

Attribute of the `componentScopedRefs` element that specifies the link to the corresponding `<component-scoped-refs>` element in `webservicesclient.xml` file.

You can edit this property in the Assembly Toolkit:

1. Open the **Web Services Client Bindings** editor.
2. Expand the **Component scoped references** section.
3. Click **Add**.
4. Select the component scoped references defined in the `webservicesclient.xml` file from the list.

- **serviceRefLink**

Attribute of the `serviceRefs` element that specifies the link to the `<service-ref-name>` in the `webservicesclient.xml` file.

You can edit this property in the Assembly Toolkit:

1. Open the **Web Services Client Bindings** editor.
2. Click the **Services References** tab.
3. Click **Add**.
4. Select the service references defined in the `webservicesclient.xml` file from the list.

- **deployedWSDLFile**

Attribute of the `serviceRefs` element is optional and permits an alternate WSDL file to be used other than that specified in the `<wsdl-file>` element of `webservicesclient.xml` file. If this attribute is specified, the alternate WSDL file must be packaged in the same module and must be compatible with the development WSDL file. The `deployedWSDLFile` property is used to supply a new WSDL file containing a different endpoint URL than the original WSDL file.

You can edit this property in the Assembly Toolkit:

1. Open the **Web Services Client Bindings** editor.
2. Select the service references or component scoped reference desired.
3. Expand the **Service reference details** section.
4. Click **Browse** on the **Deployed WSDL file** field.
5. Select the new WSDL file.
6. Click **OK**.

- **defaultMappings** element

Identifies which port should be used for a given portType when none is explicitly selected by the client. This element has the following attributes: portTypeNamespace, portTypeLocalName, portNamespace, portLocalName. These attributes identify which wsdl:port should be used for a wsdl:portType.

You can edit this property in the Assembly Toolkit:

1. Open the **Web Services Client Bindings** editor.
2. Click **Default Mappings**.
3. Click **Add**.
4. Edit the entries in the newly added row to establish a mapping between a *portType* and *port* in the WSDL file. There can only be one entry for each *portType*.
5. Click **OK**.

- **syncTimeout**

Attribute of the portQnameBindings element that specifies how long, in seconds, to wait for a response from a synchronous call.

You can edit this property in the Assembly Toolkit:

1. Create a **Port Qualified Name Bindings** for the port.
  - a. Open the **Web Services Client Bindings** editor.
  - b. Confirm that a service reference is selected in either the **Component scoped references** or **Service references** section.
  - c. Expand the **Port qualified name bindings** section.
  - d. Click **Add**. A new entry is now added to the **Port qualified name bindings** list.
2. Click the new **Port qualified name bindings** entry. The **Web Services Client Port Bindings** editor opens.
3. Expand the **Port qualified name bindings details** section.
4. Type the *namespace* of the WSDL file port you want to configure, in the **Port Namespace Link** field.
5. Type the *local\_name* of the WSDL file port you want to configure in the **Port Local Name Link** field. The name displayed in the **Port qualified name bindings** list is the local name of the WSDL file port.
6. Click **OK**.

To configure the syncTimeout property, locate the **Synchronization timeout** field and enter the desired value.

- **basicAuth**

Element of the portQnameBindings element that can be used to authenticate a service client to the service endpoint, independent of the underlying transport that includes, HTTP, HTTPS, and JMS. Set the user ID and password attributes as needed.

You can edit this property in the Assembly Toolkit:



1. Open the **Web Services Client Bindings** editor.
2. Expand the **Basic authentication** section.
3. Type the desired value in the **User ID** and **Password** fields.
4. Click **OK**.

- **sslConfig**

Element of the portQnameBindings element that specifies the Secure Sockets Layer (SSL) configuration of an HTTPS outbound request. The name attribute is the name of a SSL configuration entry or alias defined in the SSL Configuration Repertoire.

**Note:** This attribute is only used when the client is running in the WebSphere Application Server.

You can edit this property in the Assembly Toolkit:

1. Open the **Web Services Client Bindings** editor.
2. Expand the **SSL Configuration** section.
3. Type the desired value in the **Name** field.
4. Click **OK**.

The values of deployedWSDLFile and the defaultMappings of a deployed Web service can also be changed using the administrative console. Using application management, navigate to the Web module or EJB module of the Web service application and select **Web Services Client Bindings**.

### Example bindings files

The following examples demonstrate the spelling and position of the various attributes. You cannot cut and paste these examples because they do not contain the required ID attributes. If you add elements to a binding file template generated by the **WSDL2Java** command, you must confirm that each element has an ID attribute whose value is a unique string. Review the template xmi files generated by the **WSDL2Java** command for examples of ID strings.

#### Example ibm-webservices-bnd.xmi file

```
<com.ibm.etools.webservice.wsbind:WSBinding xmi:version="2.0"
  xmlns:xmi="http://www.omg.org/XMI"
  xmlns:com.ibm.etools.webservice.wsbind=
    "http://www.ibm.com/websphere/appserver/schemas/5.0.2/wsbind.xmi">
  <wsdescBindings wsDescNameLink="AddressBookService">
    <pcBindings pcNameLink="AddressBook" scope="Application"/>
  </wsdescBindings>
</com.ibm.etools.webservice.wsbind:WSBinding>
```

#### Example ibm-webservicesclient-bnd.xmi file

```
<com.ibm.etools.webservice.wscbind:ClientBinding xmi:version="2.0"
  xmlns:xmi="http://www.omg.org/XMI"
  xmlns:com.ibm.etools.webservice.wscbind=
    "http://www.ibm.com/websphere/appserver/schemas/5.0.2/wscbind.xmi">
  <componentScopedRefs componentNameLink="myComponent ref"/>
  <serviceRefs serviceRefLink="myService ref"
    deployedWSDLFile="META-INF/wsdl/alternate.wsdl">
  <defaultMappings portTypeLocalName="AddressBook"
    portTypeNamespace="http://www.com.ibm"
    portLocalName="AddressBookPort" portNamespace="http://www.com.ibm"/>
  <portQnameBindings portQnameNamespaceLink="http://www.com.ibm"
    portQnameLocalNameLink="AddressBookPort" syncTimeout="99">
```

```
<basicAuth userid="myId" password="myPassword"/>
<sslConfig name="mynode/DefaultSSLSettings"/>
</portQnameBindings>
</serviceRefs>
</com.ibm.etools.webservice.wscbind:ClientBinding>
```

## Enabling the EAR file

### Before you begin

The **endptEnabler** command is used to add router modules to your Web services-enabled application, also known as an enterprise archive (EAR) file. A router module provides an endpoint for the Web services in a particular enterprise JavaBean (EJB) Java archive (JAR) module. The **endptEnabler** command adds one or more router modules to the EAR file for each EJB JAR module within the EAR file.

Each router module supports a specific transport such as HyperText Transport Protocol (HTTP) or Java Messaging Service (JMS). If there are no EJB JAR modules in the EAR file, it is not necessary to run the **endptEnabler** command.

You are prompted for various input values for each Web services-enabled EJB JAR module in the EAR file. Typically, you should accept the defaults for each prompt. See **endptEnabler** prompts and commands for more information about **endptEnabler** command prompts.

To run the command:

### Steps for this task

1. Invoke the **endptEnabler** command from the *install\_root*\bin directory. If you are using UNIX, invoke the command from the *install\_root*/bin directory.
2. Enter the name of the EAR file, when prompted.
3. Enter various input values as they are requested by the **endptEnabler** command.

### Results

An HTTP router module is added to the EAR file for each Web services-enabled EJB JAR module contained in the EAR file. A context-root is configured for the application so the Web service can be invoked through a URL. The URL used to invoke the Web service is:

```
http://host[:port]/context-root/services/port-component-name
```

### What to do next

Deploy the EAR file into WebSphere Application Server.

### **endptEnabler** command

The **endptEnabler** enables a set of Web services within an enterprise archive (EAR) file. You can add one or more router modules to the EAR file that include a Web service-enabled EJB JAR file.

Each router module provides a Web service endpoint for a particular transport. For example, an Hyper Text Transport Protocol (HTTP) router module can be added so that the Web service can receive requests over the HTTP transport, and a Java

Messaging Service (JMS) router module can be added so that the Web service can receive requests from a JMS queue or topic.

In its interactive mode, the `endptEnabler` guides you through the required steps to enable one or more services within an application. The `endptEnabler` tool makes a backup copy of your original EAR file in the event that you need to remove or add services at a later time. If your EAR file contains a Web service-enabled EJB JAR file, you must run `endptEnabler` before the EAR file is deployed. Otherwise, you do not need to run `endptEnabler`.

### **endptEnabler usage syntax**

Invoke the `endptEnabler` from the WebSphere Application Server `bin` directory. The command syntax is as follows:

```
endptEnabler
  [-verbose|-v]
  [-quiet|-q]
  [-help|-h|-?]
  [-properties|-p <properties-filename>]
  [-transport|-t <default-transport>]
  [-enableHttpRouterSecurity]
  [<ear-filename>]
```

All parameters are optional and described as follows:

- **-verbose, -v**  
Detailed progress messages are displayed as the tool processes the EAR file. This command-line option is mapped to the `verbose` global property.
- **-quiet, -q**  
No per-module progress messages are displayed as the tool processes the EAR file. This command-line option is mapped to the `quiet` global property.
- **-help, -h, -?**  
A brief help message is displayed explaining the various options.
- **-properties, -p <properties-filename>**  
Properties from the file `<properties-filename>` are read and used to control the behavior of the tool.
- **-transport, -t <default-transport>**  
Specifies the default list of transports for which router modules should be created for each EJB JAR file contained in the EAR file. This command-line option is mapped to the `defaultTransports` global property. Examples are:  

```
-transport http (the default)
-transport jms
-t http,jms
```
- **-enableHttpRouterSecurity**  
Enables you to add a security policy for all authenticated users to protect the HTTP router module if all the EJB's are secured in the EJB JAR file. This command-line option is mapped to the `http.enableRouterSecurity` global property.
- **<ear-filename>**  
Specifies the name of the EAR file to be processed.  
If the `<ear-filename>` parameter is not entered on the command line, the interactive mode is used. In interactive mode, you are prompted for the EAR file

name, router module names and other important values as the processing occurs. The following dialog is an example of the endptEnabler interactive mode:

**Note:** In this dialog, user input is in fixed width font, and endptEnabler output is in bold.

```
endptEnabler<enter>
WSWS2004I: IBM WebSphere Application Server Release 5
WSWS2005I: Web Services Enterprise Archive Endpoint Enabler Tool.
WSWS2007I: (C) COPYRIGHT International Business Machines Corp. 1997, 2003
WSWS2006I: Please enter the name of your EAR file: AddressBook.ear
<enter>

WSWS2003I: Backing up EAR file to: AddressBook.ear~

WSWS2016I: Loading EAR file: AddressBook.ear
WSWS2017I: Found EJB Module: AddressBookEJB.jar

WSWS2029I: Enter http router name for EJB Module AddressBookEJB [AddressBookEJB_HTTPRouter.war]:
<enter>
WSWS2030I: Enter http context root for EJB Module AddressBookEJB [/AddressBookEJB]:
<enter>
WSWS2024I: Adding http router for EJB Module AddressBookEJB.jar.
WSWS2036I: Saving EAR file AddressBook.ear...
WSWS2037I: Finished saving the EAR file.
WSWS2018I: Finished processing EAR file AddressBook.ear.
```

If the <ear-filename> parameter is entered on the command line, the non-interactive mode is used. In non-interactive mode, router module names and other important values are determined from user-specified properties or default values.

### endptEnabler properties

The endptEnabler tool allows you to control its run time behavior by specifying a set of properties with the -properties command-line option. These properties fall into two categories: global and per-module. Global properties affect the overall behavior of the tool as it processes multiple EJB JAR modules within the EAR file. Per-module properties affect the processing of a particular EJB JAR module.

### Global properties

The following table describes the global properties supported by endptEnabler:

| Property name             | Description  | Default value |
|---------------------------|--|---------------|
| verbose                   | Displays detailed progress messages.   | False         |
| quiet                     | Displays only brief progress messages.   | False         |
| http.enableRouterSecurity | Enables you to add a security policy for all authenticated users to protect the HTTP router module if all the EJB's are secured in the EJB JAR file. | False         |

|                                 |  |               |
|---------------------------------|--|---------------|
| http.routerModuleNameSuffix     | Specifies the suffix used to construct default HTTP router module names. The .war extension is added by endptEnabler.  | _HTTPRouter   |
| jms.routerModuleNameSuffix      | Specifies the suffix used to construct default JMS router module names. The .jar extension is added by endptEnabler.   | _JMSRouter    |
| jms.listenerInputPortNameSuffix | Specifies the suffix used to construct default Listener Input Port names.  | _ListenerPort |
| jms.defaultDestinationType      | Specifies the default destination type to use for all JMS router modules added to the EAR file. This should be either queue or topic.  | queue         |
| defaultTransports               | Specifies the default list of transports for which router modules should be created. The list can contain the values http and jms. Multiple values are separated by a comma. Examples are: http, jms and http,jms. | http          |

### Per-module properties

The following table describes the per-module properties supported by the endptEnabler. <ejbJarName> refers to the name of an EJB JAR module within the EAR file, without the .jar extension.

| Property name                      | Description  | Default value           |
|------------------------------------|--|-------------------------|
| <ejbJarName>.transports            | Lists the transports for which router modules should be created for a particular EJB JAR file. The list can contain the values http and jms. Multiple values are separated by a comma. Examples are: http, jms and http,jms. | http                    |
| <ejbJarName>.http.skip             | Specifies the flag which bypasses the addition of an HTTP router module even if it would otherwise be added (based on other properties). Valid values are true and false.  | False                   |
| <ejbJarName>.http.routerModuleName | Specifies the name of the HTTP router module for a particular EJB JAR file.  | <ejbJarName>_HTTPRouter |

|  |   |  |
|--|---|--|
| <code>&lt;ejbJarName&gt;.http.contextRoot</code>                   | Specifies the context root associated with the HTTP router module for a particular EJB JAR file.  | <code>/&lt;ejbJarName&gt;</code>             |
| <code>&lt;ejbJarName&gt;.jms.skip</code>                           | Specifies the Flag which bypasses the addition of an HTTP router module even if it would otherwise be added (based on other properties). Valid values are true and false. | false  |
| <code>&lt;ejbJarName&gt;.jms.routerModuleName</code>               | Specifies the name of the JMS router module for a particular EJB JAR file.  | <code>&lt;ejbJarName&gt;_JMSRouter</code>    |
| <code>&lt;ejbJarName&gt;.jms.listenerInputPortName</code>          | Specifies the name of the Listener Input Port to be associated with the JMS router module.  | <code>&lt;ejbJarName&gt;_ListenerPort</code> |
| <code>&lt;ejbJarName&gt;.ejbJarName&gt;.jms.destinationType</code> | Specifies the JMS destination type associated with the JMS router. Valid values are queue and topic.  | queue  |

### Properties example

Suppose an EAR file contains an EJB JAR file named, StockQuoteEJB.jar that contains Web services. The following set of properties might be used to control the endptEnabler runtime behavior as it processes the EAR file:

```
StockQuoteEJB.transports=http,jms

StockQuoteEJB.http.routerModuleName=StockQuoteEJB_HTTP

StockQuoteEJB.http.contextRoot=/StockQuote

StockQuoteEJB.jms.routerModuleName=StockQuoteEJB_JMS

StockQuoteEJB.jms.listenerInputPortName=StockQuote_LP

StockQuoteEJB.jms.destinationType=queue
```

### endptEnabler examples

The following commands are examples of how the endptEnabler can be used:

```
endptEnabler MyApp.ear

endptEnabler -t jms,http MyApp.ear

endptEnabler -v -properties MyApp.props MyApp.ear

endptEnabler -q -t jms MyApp.ear
```

---

## Deploying Web services

### Before you begin

You need an enterprise application, also known as an enterprise archive (EAR) file, that has been configured and enabled for Web services. You can use either the administrative console or the **wsadmin** scripting interface to deploy an EAR file.

If you are installing an application containing Web services by using the **wsadmin** command, specify the deployment option. If you are installing an application containing Web services by using the administrative console, select **Deploy WebServices** during step 1 of the Install New Application wizard.

**Note:**

If the Web services in the application is previously deployed with the **wsdeploy** command, it is not necessary to specify Web services deployment during installation.

Use the following steps to deploy the EAR file with the **wsadmin** command:

Steps for this task

1. Start *install\_root*\bin\wsadmin from a command prompt.  
If you are using UNIX start *install\_root*/bin/wsadmin.
2. Enter the **\$AdminApp install EARfile "-usedefaultbindings -deployejb -deployws"** command at the **wsadmin** prompt.

Results

The Web service is installed into the application server.

What to do next

Secure Web services.

## **wsdeploy command**

The **wsdeploy** command line tool adds Websphere product-specific Web services deployment classes to a Web services for J2EE (JSR-109) compatible enterprise application enterprise archive (EAR) file or an application client Java archive (JAR) file. These classes include:

- Stubs
- Serializers and deserializers
- Implementations of service interfaces

This deployment step must be performed at least once, and can be performed more than once. Deployment can be performed separately using the **wsdeploy** command, the Assembly Toolkit (<http://www-3.ibm.com/software/webservers/appserv/was/support/>), or when the application is installed. When using the **wsadmin** command for installation, specify the **-deployws** option. When using the administrative console for installation, select the **Deploy Web services** check box. When using the Assembly Toolkit, **Right-click** the module and select **Deploy Web Services** from the pop-up menu.

The **wsdeploy** command operates as follows:

- Each module in the enterprise application or JAR file is examined

- If the module contains Web services implementations, indicated by the presence of the `webservices.xml` deployment descriptor, the associated Web Services Description Language (WSDL) files are located and the **WSDL2Java** command is run with the role `deploy-server`.
- If the module contains Web services clients, indicated by the presence of the `webservicesclient.xml` deployment descriptor, the associated WSDL files are located and the **WSDL2Java** command is run with the role `deploy-client`.
- The files generated by the **WSDL2Java** command are compiled and repackaged.

See `WSDL2Java` command for more information about the files that are generated for deployment.

When the generated files are compiled, they can reference application-specific classes outside the EAR or JAR file if the EAR or JAR file is not self-contained. In this case, use either the `-jardir` or `-cp` option to specify additional JAR or zip files to be added to `CLASSPATH` when the generated files are compiled.

### **wsdeploy command syntax**

The command syntax is as follows:

```
wsdeploy Input_filename Output_filename [options]
```

#### **Required options:**

- ***Input\_filename***  
Specifies the path to the EAR or JAR file to be deployed.
- ***Output\_filename***  
Specifies the path of the deployed EAR or JAR file. If `output_filename` already exists, it is silently overwritten. The `output_filename` can be the same as the `input_filename`.

#### **Other options:**

- ***-jardir directory***  
Specifies a directory containing JAR or zip files. All JAR and zip files in this directory are added to the `CLASSPATH` used to compile the generated files. This option can be specified zero or more times.
- ***-cp entries***  
Specifies entries to be added to `CLASSPATH` when the generated classes are compiled. Multiple entries are separated the same as they would be in the `CLASSPATH` environment variable, with a semicolon on Windows platforms and a colon for UNIX platforms.
- ***-codegen***  
Specifies that deployment code is to be generated, but not compiled. This option implicitly specifies the `-keep` option.
- ***-debug***  
Includes debugging information when compiling, that is, use `javac -g` to compile.
- ***-help***  
Displays a help message and exit.
- ***-ignoreerrors***  
Do not stop deployment if validation or compilation errors are encountered.
- ***-keep***



Do not delete working directories containing generated classes. A message is displayed indicating the name of the working directory that is retained.

- **-novalidate**

Do not validate the Web services deployment descriptors in the input file.

- **-trace**

Displays processing information, including the names of the generated files.

### Example

```
wsdeploy x.ear x_deployed.ear -trace -keep
Processing web service module x_client.jar.
Keeping directory: f:\temp\Base53383.tmp for module: x_client.jar.
Parsing XML file:f:\temp\Base53383.tmp\WarDeploy.wsdl
Generating f:\temp\Base53383.tmp\generatedSource\com\test\WarDeploy.java
Generating f:\temp\Base53383.tmp\generatedSource\com\test\WarDeployLocator.java
Generating f:\temp\Base53383.tmp\generatedSource\com\test\HelloWsBindingStub.java
Compiling f:\temp\Base53383.tmp\generatedSource\com\test\WarDeploy.java.
Compiling f:\temp\Base53383.tmp\generatedSource\com\test\WarDeployLocator.java.
Compiling f:\temp\Base53383.tmp\generatedSource\com\test\HelloWsBindingStub.java.
Done processing module x_client.jar.
```

### Messages

- Flag *-f* is not valid

Option *f* was not recognized as being a valid option.

- Flag *-c* is ambiguous

Options may be abbreviated, but the abbreviation must be unique. In this case, the **wsdeploy** command can not determine which option was intended.

- Flag *-c* is missing parameter *-p*

A required parameter for an option was omitted.

- Missing *p* parameter

A required option was omitted.

---

## Using Java Messaging Service to transport Web services requests

### Before you begin

WebSphere Application Server offers support for using a Java Messaging Service (JMS) transport layer, in addition to the existing HTTP transport. Using JMS transport allows your Web service clients and servers to communicate through JMS queues and topics instead of HTTP connections. One-way and synchronous two-way requests are supported.

**Note:** A Web service must be implemented as an enterprise JavaBean (EJB) to be accessed through the JMS transport.

The benefits of using JMS as an alternative to HTTP, include:

- Request and response messages are sent through reliable messaging.
- One-way requests allow clients and servers to be more loosely-coupled. For example, the server does not have to be active when the client sends the one-way request.
- One-way requests can be sent to multiple servers simultaneously through the use of a topic.

To use JMS as a transport for Web services requests:

### Steps for this task

1. Add a JMS binding and a Simple Object Access Protocol (SOAP) address to the Web Services Description Language (WSDL) file.

The WSDL file of a Web service must include a JMS binding and a SOAP address, which specifies a JMS endpoint URL string, in order to be accessible on the JMS transport. A JMS binding is a `wsdl:binding` element containing a `wsdlsoap:binding` element whose `transport` attribute ends in `soap/jms`, rather than the typical `soap/http` value.

In addition to the JMS binding, a `wsdl:port` element referencing the JMS binding must be included in the `wsdl:service` element within the WSDL file. The `wsdl:port` element should contain a `wsdlsoap:address` element whose `location` attribute specifies a JMS endpoint URL string.

**Note:** The specification of the actual JMS endpoint URL string can be deferred until you publish the WSDL file. When you develop the Web service, a placeholder such as `file:/unspecified_location` can be used for the endpoint URL string.

2. Decide on the names and types of JMS objects that your application uses

Before your application can be installed, you need to:

- a. Decide whether your Web service receives its requests from a queue or a topic.
- b. Decide whether to use a secure destination (queue or topic) or a nonsecure destination.
- c. Decide on the names for your destination, connection factory and listener port.

The following list provides examples of the names that might be used for the mythical StockQuote Web service:

- **Queue:** StockQuote\_Q (JNDI name: `jms/StockQuote_Q`)
- **Connection factory:** StockQuote\_CF (JNDI name: `jms/StockQuote_CF`)
- **Listener port:** StockQuoteEJB\_ListenerPort

3. Define the JMS administered objects.

Once you have decided on the names and types of the JMS objects, use the administrative console or the **wsadmin** scripting interface to define the JMS objects.

4. Add the JMS endpoints to your EAR file using the **endptEnabler** command tool.

You must run the **endptEnabler** command to add a JMS endpoint or router module for each Web service-enabled EJB JAR file contained in the EAR file. By default, the **endptEnabler** command adds only HTTP endpoints, but the `-transport jms` option can be used to request the addition of JMS endpoints.

5. Deploy the Web services application.

During the install process you are prompted for two types of information for each Web service-enabled EJB JAR contained in your EAR file:

- The Java Naming and Directory Interface (JNDI) name of the connection factory to be used by the EJB JAR file message driven bean (MDB) listener for sending reply messages.

If your Web service contains two-way operations, the MDB listener, defined inside the JMS endpoint added by **endptEnabler** command, needs to access a queue connection factory in order to add a reply message to the reply queue.

The MDB listener uses a resource environment reference of `java:comp/env/jms/WebServicesReplyQCF`. Therefore, during the application

install process, you must provide the actual JNDI name of the queue connection factory that should be used by the MDB listener for that Web service. You might want to use the same connection factory that you defined for use by clients in step 2.

- The name of the listener port to be used by the MDB listener.  
A listener port is an object used to associate a JMS connection factory with a JMS destination (queue or topic). When deployed, an MDB is configured with the correct listener port so that messages from the desired queue or topic are properly delivered to the MDB. During deployment, you can modify the name of the listener port associated with each MDB listener. The listener port name contained in the input EAR file displays as a default value. If you specify the correct listener port name to the **endptEnabler** command, perhaps through the use of properties, during step 3, you can accept the default value. Otherwise, enter the correct listener port name.

**Hint:** By default, the **endptEnabler** command produces listener port names of the form `<ejb-jar-name>_ListenerPort`. To simplify this step, define the listener ports that follow this naming convention during step 2.

#### 6. Publish the WSDL file.

In this step, you enter the JMS endpoint URL string to use for each Web service-enabled EJB JAR file belonging to the application. The JMS endpoint URLs are then written to the published WSDL files for use by clients.

For example, suppose that an application called `StockQuoteService` contains an EJB JAR file named `StockQuoteEJB`, which contains one or more Web services accessible on the JMS transport. Suppose that, in step 2, you defined a queue with the JNDI name `jms/StockQuote_Q` and a connection factory with the JNDI name `jms/StockQuote_CF` to be used by your application. In this example, you would specify the following string as the JMS URL prefix within the **Publish WSDL** user interface:

```
jms:/queue?destination=jms/StockQuote_Q&connectionFactory=jms/StockQuote_CF
```

The WSDL publisher uses this partial URL string to produce the actual JMS URL for each port component defined in the EJB JAR file. The published WSDL file can be used by clients needing to invoke the Web service.

## Java Messaging Service endpoint URL syntax

A Java Messaging Service (JMS) endpoint URL is used to access a Web service with the JMS transport. This URL specifies the JMS destination and connection factory, as well as the port component name for the Web service request. This is similar to the HTTP endpoint URL, which specifies the host and port, as well as the context root and port component name.

A JMS endpoint URL has the following general form:

```
jms:[queue|topic]?<property>=<value>&<property>=<value>&...
```

The URL consists of the transport type, `jms:`, followed by either `/queue` or `/topic` to indicate the JMS destination type, followed by the query string containing a list of property and value pairs used to specify the JMS endpoint information.

The properties supported in the URL string are described as follows:

### Destination-related properties (required)

| Property name | Description |
|---------------|-------------|
|---------------|-------------|

|                   |  |
|-------------------|--|
| destination       | Specifies the Java Naming and Directory Interface (JNDI) name of the destination queue or topic. |
| connectionFactory | Specifies the JNDI name of the connection factory.   |
| targetService     | Specifies the name of the port component to which the request is dispatched.                     |

### JNDI-related properties (optional)

| Property name         | Description   |
|-----------------------|---|
| initialContextFactory | Specifies the name of the initial context factory to use which is mapped to the java.naming.factory.initial property. |
| jndiProviderURL       | Specifies the JNDI provider URL which is mapped to the java.naming.provider.url property.                             |

### JMS-related properties (optional)

| Property name | Description   |
|---------------|---|
| deliveryMode  | Indicates whether the request message should be persistent or not. The valid values are 1 for nonpersistent and 2 for persistent. The default value is 1. |
| timeToLive    | Specifies the lifetime, in milliseconds, of the request message. A value of 0 indicates an infinite lifetime.   |
| priority      | Specifies the JMS priority associated with the request message. Valid values are between 0 to 9. The default value is 4.                                  |

The required properties, destination, connectionFactory, and targetService, must appear in the JMS endpoint URL string. The rest of the properties are optional.

You can set any of the properties on the client Stub object. This means that the various properties can be specified by including them as part of the endpoint URL or they can be set programmatically by the client on the Stub object. Properties specified on the client Stub object take precedence over properties specified as part of a JMS endpoint URL string.

---

## Securing Web services

### Before you begin

Web services security for WebSphere Application Server, Version 5.0.2 and above is based on standards included in the Web services security (WS-Security) specification. Web services security is a message-level standard, based on securing Simple Object Access Protocol (SOAP) messages through XML digital signature, confidentiality through XML encryption and credential propagation through security tokens.

Use the deprecated Securing Apache SOAP services topics if you are still using Apache SOAP 2.3.

The movement to streamline secured Web services leads to a standards-based architecture that is comprehensive, yet flexible enough to meet the Web services security needs of real businesses. These standards are a set of Web service security specifications that address how to provide protection for messages exchanged in a Web service environment. Web services security defines the core facilities for protecting the integrity and confidentiality of a message and provides mechanisms for associating security-related claims with the message.

In the real world of e-business, qualities of services are required to provide security, reliable messaging, and management for each layer of the Web services stack. To secure Web services, you must consider a broad set of security requirements, including authentication, authorization, privacy, trust, integrity, confidentiality, secure communications channels, federation, delegation, and auditing across a spectrum of application and business topologies. One of the key requirements for the security model in today's business environment is the ability to interoperate between formerly incompatible security technologies (such as public key infrastructure, Kerberos and so on.) in heterogeneous environments (such as .NET and J2EE). The complete Web services security protocol stack and technology roadmap is described in "Security in a Web Services World: A Proposed Architecture and Roadmap", which is available at the following address: <http://www.ibm.com/developerworks/webservices/library/ws-secmap/> (<http://www.ibm.com/developerworks/webservices/library/ws-secmap/>).

The Web services security specification, which is available at <http://www.ibm.com/developerworks/library/ws-secure/> (<http://www.ibm.com/developerworks/library/ws-secure/>), proposes a standard set of Simple Object Access Protocol (SOAP) extensions that you can use to build secure Web services. These standards confirm integrity and confidentiality, which are generally provided with digital signature and encryption technologies. In addition, Web services security provides a general purpose mechanism for associating security tokens with messages. A typical example of the security token is a user name and password token, in which a user name and password are included as text. Web services security defines how to encode binary security tokens such as X.509 certificates and Kerberos tickets.

For an explanation of Web services security and for instructions on how to configure your WebSphere Application Server, see Web services security for WebSphere Application Server, Version 5.0.2.

## Configuring client-side transport level security

### Before you begin

The server-side, or service endpoint, transport level security is based on the Secured Sockets Layer (SSL) configuration of the WebSphere Application Server Web container. Review *Configuring Secure Sockets Layer* for more information.

To configure the client-side transport level security:

### Steps for this task

1. Create a SSL repertoire configuration entry for an existing service endpoint acting as a service client.

OR, you can use an existing configuration alias from the SSL configuration repertoire using the Application Assembly Tool (AAT).

2. Define the attribute `sslConfig` with the value of the alias name in the `ibm-webservicesclient-bnd.xmi` file.

For example:

```
<sslConfig name="default/DefaultSSLSettings"/>
```

**Note:** If the attribute is not defined, the default SSL setting is used for JSSE.

3. Locate the property file `<install_root>\properties\sas.client.props` for the service clients or create a new property file that includes:

```
com.ibm.ssl.protocol
com.ibm.ssl.keyStoreType
com.ibm.ssl.keyStore
com.ibm.ssl.keyStorePassword
com.ibm.ssl.trustStoreType
com.ibm.ssl.trustStore
com.ibm.ssl.trustStorePassword
```

4. Set the system property, `com.ibm.webservices.sslConfigURL` to the property file.

For example:

```
Dcom.ibm.webservices.sslConfigURL=<installation_root>\properties\sas.client.props
```

**Note:** If the property `sslConfigURL` is not defined, the default SSL setting is used for JSSE.

5. (Optional) Set the system properties of an unmanaged service client by using the `-D` option of the Java command or by calling the `System.setProperty` (`propertyName, "<samp> propertyValue "`) with the following properties:

```
java.protocol.handler.pkgs
java.net.ssl.keyStore
javax.net.ssl.keyStorePassword
javax.net.ssl.trustStore
javax.net.ssl.trustStorePassword
```

See [Using Java Secure Socket Extension \(JSSE\) and Java Cryptography Extension \(JCA\) with servlets and enterprise bean files](#) for more information about customizing the JSSE.

6. Access the service endpoints in a Federal Information Processing Standard (FIPS)-enabled WebSphere Application Server.
  - a. Check for the required properties defined in the WebSphere Application Server security documentation.

7. (Optional) Redirect the Simple Object Access Protocol (SOAP) request from a client to service endpoint to be over HTTPS.

Complete this step if a transport guarantee of **CONFIDENTIAL** or **INTEGRAL** is configured for a secured Web application. To redirect the request:

- a. Specify a system property, `com.ibm.ws.webservices.HttpRedirectEnabled`, to **true** for the entire Java Virtual Machine (JVM) or set the property, `com.ibm.wsspi.webservices.Constants.HTTP_REDIRECT_ENABLED`, to **true**, in the stub or call instance, before the method is invoked.

## Transport level security

*Transport level security* is based on Secured Sockets Layer (SSL) or Transport Layer Security (TLS) that runs beneath the HTTP protocol.

SSL and TLS provide security features including authentication, data protection, and cryptographic token support for secure HTTP connections. To run with HTTPS, the service endpoint address must be in the form of `https://`.

The integrity and confidentiality of transport data, including Simple Object Access Protocol (SOAP) messages and HTTP basic authentication, is confirmed when you use SSL and TLS. See Secured Sockets Layer for more information. Web services applications can also use Federal Information Processing Standard (FIPS) approved ciphers for more secure TLS connections.

WebSphere Application Server uses Java Secure Sockets Extension (JSSE) to support SSL and TLS.

## Configuring HTTP basic authentication

### Before you begin

You can configure the username and password for HTTP basic authentication with the `ibm-webservicesclient-bnd.xmi` deployment descriptor, or using the properties mechanism for the configuration of a stub or call instance at run time.

To configure the HTTP basic authentication:

### Steps for this task

1. Using the `ibm-webservicesclient-bnd.xmi` deployment descriptor: Specify the attribute `basicAuth` for each `port QNameBindings` of each `serviceRef`. For example:

```
<basicAuth userid="myID" password="myPassword"\>
```

2. Using the properties mechanism: Change the following properties:

```
javax.xml.rpc.Call.USERNAME_PROPERTY  
javax.xml.rpc.Call.PASSWORD_PROPERTY  
javax.xml.rpc.Stub.USERNAME_PROPERTY  
javax.xml.rpc.Stub.PASSWORD_PROPERTY
```

**Note:** The values set by the properties mechanism take precedence over the values defined by the `ibm-webservicesclient-bnd.xmi` deployment descriptor. You can also configure Proxy data using the properties mechanism described by using the following properties to configure your Web services application:

For HTTP:

```
com.ibm.wsspi.webservices.HTTP_PROXYHOST_PROPERTY  
com.ibm.wsspi.webservices.HTTP_PROXYPORT_PROPERTY  
com.ibm.wsspi.webservices.HTTP_PROXYUSER_PROPERTY  
com.ibm.wsspi.webservices.HTTP_PROXYPASSWORD_PROPERTY
```

For HTTPS:

```
com.ibm.wsspi.webservices.HTTPS_PROXYHOST_PROPERTY  
com.ibm.wsspi.webservices.HTTPS_PROXYPORT_PROPERTY  
com.ibm.wsspi.webservices.HTTPS_PROXYUSER_PROPERTY  
com.ibm.wsspi.webservices.HTTPS_PROXYPASSWORD_PROPERTY
```

### **HTTP basic authentication**

*HTTP basic authentication* uses a username and password to authenticate a service client to a secure endpoint.

WebSphere Application Server can have several resources, including Web services, protected by a J2EE security model.

A simple way to provide authentication data for the service client is to authenticate to the protected service endpoint to the HTTP basic authentication. The basic authentication is located in the HTTP header that carries the Simple Object Access Protocol (SOAP) request. When the application server receives the HTTP request, the username and password are retrieved and verified using the authentication mechanism specific to the server.

Although the basic authentication data is base64-encoded, it is recommended that the data is sent over HTTPS. The integrity and confidentiality of the data can be protected by the Secured Sockets Layer (SSL) protocol.

In some cases, a firewall is present using the PASS-THRU HTTP proxy server. The HTTP proxy server forwards the basic authentication data into the J2EE application server. The proxy server can also be protected. Applications can specify the proxy data by setting properties in a stub object.

## Web Services: Default bindings for the Web Services Security collection

Use this page to manage the default bindings for trust anchors, the collection certificate store, key locators, trusted ID evaluators, and login mappings. The default binding configuration provides a central location where reusable binding information is defined. The application binding file can reference the information contained in the default binding configuration.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Web Services: Default bindings for Web Services Security**.

Read the Web Services documentation in the InfoCenter before you begin defining the default bindings for Web Services Security.

Define the server bindings file by clicking **Enterprise Applications > *application\_name***. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Server Security Bindings**.

Define the client bindings file by clicking **Enterprise Applications > *application\_name***. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Client Security Bindings**.

### Trust Anchors

Specifies a list of key store objects that contain the trusted root certificates, self-signed or issued by a certificate authority (CA).

The certificate authority authenticates a user and issues a certificate. After the certificate is issued, the key store objects, which contain these certificates, use the certificate for certificate path or certificate chain validation of incoming X.509-formatted security tokens.

### Collection Certificate Store

Specifies a list of the untrusted, intermediate certificate files.

The collection certificate store contains a chain of untrusted, intermediate certificates. The CertPath API attempts to validate these certificates, which are based on the trust anchor.



## Key Locators

Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

## Trusted ID Evaluators

Specifies a list of trusted ID evaluators that determine whether to trust the identity-asserting authority or message sender.

The trusted ID evaluators are used to authenticate additional identities from one server to another server. For example, a client sends the identity of user A to server 1 for authentication. Server 1 calls downstream to server 2, asserts the identity of user A, and includes the user ID and password of server 1. Server 2 attempts to establish trust with server 1 by authenticating its user ID and password and checking the trust based on the TrustedIDEvaluator implementation. If the authentication process and the trust check are successful, server 2 trusts that server 1 authenticated user A and a credential is created for user A on server 2 to invoke the request.

## Login Mappings

Specifies a list of configurations for validating tokens within incoming messages.

Login mappings map the authentication method to the Java Authentication and Authorization Service (JAAS) configuration.

To configure JAAS, use the administrative console and click **Security > JAAS Configuration**.

## Trust Anchors collection

Use this page to view a list of key store objects that contain trusted root certificates. These objects are used for certificate path validation of incoming X.509-formatted security tokens. Key store objects within trust anchors contain trusted root certificates used by the CertPath API to validate the trust of a certificate chain.

To create the key store file, use the key tool located in the `%install_dir%\java\jre\bin\keytool` directory.

To view this administrative console page, click **Servers > Application Servers > server\_name**. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

Click **New** to create a new trust anchor.

Click **Delete** to a delete trust anchor.

Click **Update runtime** to update the Web services security run time with the default binding information, contained in the `ws-security.xml` file, that was previously saved. However, prior to clicking **Update runtime**, you must save your changes by clicking **Save** at the top of the Administrative console. When you click **Save**, you are returned to the Administrative console home panel. To update the run time, return to the Trust Anchors collection panel and click **Update runtime**. When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.

**Trust Anchor Name:** Specifies the unique name used to identify the trust anchor.

**Key Store Path:** Specifies the location of the key store file that contains the trust anchors.

**Key Store Type:** Specifies the type of key store.

The value for this field is either **JKS** or **JCEKS**.

### Trust Anchor configuration settings

Use this page to configure a trust anchor. Trust anchors point to key stores that contain trusted root or self-signed certificates. This page enables you to specify a name for the trust anchor and the information needed to access a key store. The name is used by the application binding to reference a pre-defined trust anchor definition in the binding file or the default binding.

To view this administrative console page, click **Servers > Application Servers > server\_name**. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors > New**.

**Trust Anchor Name:** Specifies the unique name used by the application binding to reference a pre-defined trust anchor definition in the default binding.

**Key Store Password:** Specifies the password needed to access the key store file.

**Key Store Path:** Specifies the location of the key store file.

Use `${USER_INSTALL_ROOT}` as this path expands to the WebSphere Application Server path on your machine.

**Key Store Type:** Specifies the type of key store file.

The value in this field is either **JKS** or **JCEKS**. The following is an explanation of these two options:

**JKS** Specify this option if you are not using Java Cryptography Extensions (JCE).

**JCEKS** Specify this option if you are using Java Cryptography Extensions. Although the JCEKS key store format is more secure, it decreases performance.

|           |            |
|-----------|------------|
| Data type | String     |
| Default   | JKS        |
| Range     | JKS, JCEKS |

### Collection Certificate Store collection

Use this page to view a list of certificate stores containing untrusted, intermediary certificate files awaiting validation. Validation might consist of checking to see if the certificate is on a certificate revocation list (CRL), checking that the certificate has not expired, and checking that the certificate was issued by a trusted signer.

To view this administrative console page, click **Servers > Application Servers server\_name**. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

Click **New** to specify a store name and provider for a new certificate store.

Click **Delete** to delete a certificate store.

Click **Update runtime** to update the Web services security run time with the default binding information, contained in the `ws-security.xml` file, that was previously saved. However, prior to clicking **Update runtime**, you must save your changes by clicking **Save** at the top of the Administrative console. When you click **Save**, you are returned to the Administrative console home panel. To update the run time, return to the Collection Certificate Store collection panel and click **Update runtime**. When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.

**Certificate Store Name:** Specifies the name of the certificate store.

**Certificate Store Provider:** Specifies the provider of the certificate store.

### Collection Certificate Store configuration settings

Use this page to specify the name and provider of a certificate store.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store > New**.

**Certificate Store Name:** Specifies the name for the certificate store. The application binding uses the certificate store name to reference a pre-defined binding.

**Certificate Store Provider:** Specifies the provider for the certificate store implementation.

|           |              |
|-----------|--------------|
| Data type | String       |
| Default   | IBM CertPath |

### X.509 certificates collection

Use this page to view a list of X.509 certificates.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**. On the Collection Certificate Store page, click **X.509 Certificates** under Additional Properties.

Click **New** to create a new path to an X.509 certificate.

Click **Delete** to delete a path to a X.509 certificate.

**X509 Certificate Path:** Specifies the location of the X.509 certificate.

### X.509 Certificate configuration settings

Use this page to specify the location of your X.509 certificates.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**. On the Collection Certificate Store page, under Additional Properties, click **X.509 Certificates > New**.

**X509 Certificate Path:** Specifies the location of the X.509 certificate.

## Key Locator collection

Use this page to view a list of available key locators. Key locators identify the keys needed for digital signature and encryption. A key locator must implement the `com.ibm.wsspi.wssecurity.config.KeyLocator` interface. The two default implementations are:

`com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator` and  
`com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator`.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

Click **New** to create a key locator.

Click **Delete** to delete a key locator.

Click **Update runtime** to update the Web services security run time with the default binding information, contained in the `ws-security.xml` file, that was previously saved. However, prior to clicking **Update runtime**, you must save your changes by clicking **Save** at the top of the Administrative console. When you click **Save**, you are returned to the Administrative console home panel. To update the run time, return to the Key Locator collection panel and click **Update runtime**. When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.

**Note:** Once you define key locators, click the key locator name to specify additional properties and keys under **Additional Properties**.

**Key Locator Name:** Specifies the unique name of the key locator.

**Key Locator Classname:** Specifies the class name of the key locator in the key store file.

## Key Locator configuration settings

Use this page to specify the settings for key locators.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > New**.

**Key Locator Name:** Specifies the name of the key locator.

|           |        |
|-----------|--------|
| Data type | String |
| Units     | N/A    |
| Default   | N/A    |
| Range     | N/A    |

**Key Locator Classname:** Specifies the name for the key locator class implementation.

WebSphere Application Server has the following default key locator class implementations:

### **`com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator`**

This class, used by the response sender, maps an authenticated identity to a key. If encryption is used, this class is used to locate a key to encrypt the

response message. The `com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator` class has the capability to map an authenticated identity from the invocation credential of the current thread to a key that is used to encrypt the message. If an authenticated identity is present on the current thread, the class maps the ID to the mapped name. For example, `user1` is mapped to `mappedName_1`. Otherwise, `name="default"`. When a matching key is not found, the authenticated identity is mapped to the default key specified in the binding file.

**com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator**

This class, used by the response receiver, request sender, and request receiver, maps a name to an alias. Encryption uses this class to obtain a key to encrypt a message and digital signature uses this class to obtain a key to sign a message. The `com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator` class maps a logical name to a key alias in the key store file. For example, key `#105115176771` is mapped to `CN=Alice, O=IBM, c=US`.

|           |   |
|-----------|---|
| Data type | String  |
| Defaults  | <code>com.ibm.wsspi.wssecurity.config.KeyStoreKeyLocator</code> ,<br><code>com.ibm.wsspi.wssecurity.config.WSIdKeyStoreMapKeyLocator</code> |

**Key Store Password:** Specifies the password used to access the key store file.

**Key Store Path:** Specifies the location of the key store file.

Use `${USER_INSTALL_ROOT}` as this path expands to the WebSphere Application Server path on your machine.

**Key Store Type:** Specifies the type of the key store file.

The value for this field is either **JKS** or **JCEKS**. The following is an explanation of these two options:

**JKS** Use this option if you are not using Java Cryptography Extensions (JCE).

**JCEKS** Use this option if you are using Java Cryptography Extensions.

|         |            |
|---------|------------|
| Default | JKS        |
| Range   | JKS, JCEKS |

**Key collection**

Use this page to view a list of logical names that are mapped to a key alias in the key store file.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > *key\_locator\_name***. Under Additional Properties, click **Keys**.

Click **New** to create a new key object in the key store file.

Click **Delete** to delete a mapping of a key object within the key store file.

**Key Name:** Specifies the name of the key object found in the key store file.

**Key Alias:** Specifies an alias for the key object.

The alias is used when the key locator searches for the key objects in the key store.

### Key configuration settings

Use this page to define a mapping of a logical name to a key alias in a key store file.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > *key\_locator\_name***. Under Additional Properties, click **Keys > New**.

**Key Name:** Specifies the name of the key object. This name is used by the key locator to find the key within the key store file.

**Key Alias:** Specifies the alias for the key object contained in the key store file.

**Key Password:** Specifies the password needed to access the key object within the key store file.

### Trusted ID Evaluator collection

Use this page to view a list of trusted identity (ID) evaluators. The trusted ID evaluator determines whether to trust the identity-asserting authority. Once the ID is trusted, the WebSphere Application Server issues the proper credentials, which are used in a downstream call for invoking resources. The trusted ID evaluator implements the `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluator` interface.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Services, click **Web Services: Default bindings for Web Services Security > Trusted ID Evaluators**.

Click **New** to create a trusted ID evaluator.

Click **Delete** to delete a trusted ID evaluator.

Click **Update runtime** to update the Web services security run time with the default binding information, contained in the `ws-security.xml` file, that was previously saved. However, prior to clicking **Update runtime**, you must save your changes by clicking **Save** at the top of the Administrative console. When you click **Save**, you are returned to the Administrative console home panel. To update the run time, return to the Trusted ID Evaluator collection panel and click **Update runtime**. When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.

**Trusted ID Evaluator Name:** Specifies the unique name of the trusted ID evaluator.

**Trusted ID Evaluator Classname:** Specifies the class name of the trusted ID evaluator.

### Trusted ID Evaluator configuration settings

Use this page to configure trust identity (ID) evaluators.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust ID Evaluators > New**.

You must specify the name and value properties for the default trusted ID evaluator to create the trusted ID list for evaluation.

**Trusted ID Evaluator Name:** Specifies the unique name used by the application binding to refer to a trusted identity (ID) evaluator defined in the default binding.

You must specify the trusted ID evaluator name in the form, `trustedId_n`, where `_n` is an integer from 0 to n.

**Trusted ID Evaluator Class Name:** Specifies the class name of the trusted ID evaluator.

Default      `com.ibm.wsspi.wssecurity.id.TrustedIDEvaluatorImpl`

## Login Mappings collection

Use this page to view a list of configurations for validating security tokens within incoming messages. Login mappings map an authentication method to a Java Authentication and Authorization Service (JAAS) login configuration to validate the security token. There are four authentication methods pre-defined in the WebSphere Application Server: BasicAuth, Signature, IDAssertion, and LTPA

To view this administrative console page, click **Server > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings**.

Click **New** to create a login mapping.

Click **Delete** to delete a login mapping.

Click **Update runtime** to update the Web services security run time with the default binding information, contained in the `ws-security.xml` file, that was previously saved. However, prior to clicking **Update runtime**, you must save your changes by clicking **Save** at the top of the Administrative console. When you click **Save**, you are returned to the Administrative console home panel. To update the run time, return to the Login Mappings collection panel and click **Update runtime**. When you click **Update runtime**, the configuration changes made to the other Web services also are updated in the Web services security run time.

**Authentication Method:** Specifies the authentication method used for validating the security tokens.

The following authentication methods are available:

### Basic Auth

Basic authentication includes both a user name and password in the security token. The information in the token is authenticated by the receiving server and used to create a credential.

### Signature

Used when the authentication method is signature where an X.509 certificate is sent as a security token. For Lightweight Directory Access Protocol (LDAP) registries, the distinguished name (DN) is mapped to a credential, which is based on the LDAP certificate filter settings. For Local

OS registries, the first attribute of the certificate, usually the common name (CN) is mapped directly to a user ID in the registry.

#### **IDAssertion**

Identity assertion maps a trusted identity (ID) to a WebSphere credential. This authentication method only includes a user name in the security token. An additional token is included in the message for trust purposes. Once the additional token is trusted, the IDAssertion token user name is mapped to a credential. Used when the authentication method is IDAssertion.

**LTPA** Light-weight Third Party Authentication validates an LTPA token.

**JAAS Configuration Name:** Specifies the name of the Java Authentication and Authorization Service (JAAS) configuration.

**Callback Handler Factory Class Name:** Specifies the name of the factory for the CallbackHandler class.

#### **Login Mapping configuration settings**

Use this page to specify the Java Authentication and Authorization Service (JAAS) login configuration settings used to validate security tokens within incoming messages.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings > New**.

**Authentication Method:** Specifies the method of authentication.

You can use any string, but the string must match the element in the service-level configuration. The following three words are reserved and have special meanings: "BasicAuth": uses both a user name and a password; "IDAssertion": uses only a user name, but requires that additional trust is established on the receiving server using a TrustedIDEvaluator; and "Signature": uses the distinguished name (DN) of the signer.

**JAAS Configuration Name:** Specifies the name of the Java Authentication and Authorization Service (JAAS) configuration.

Specify your JAAS configurations using the Administrative Console by clicking **Security > JAAS Configuration > Application**.

**Callback Handler Factory Class Name:** Specifies the name of the factory for the CallbackHandler class.

Use must implement the `com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory` class in this field.

Default `com.ibm.wsspi.wssecurity.auth.callback.CallbackHandlerFactory`

**Token Type URI:** Specifies the namespace URI, which denotes the type of the security token accepted.

If binary security tokens are accepted, the value should denote the ValueType attribute in element. The ValueType element identifies the type of security token



and its namespace. If XML tokens are accepted, the value should denote the top-level element name of the XML token.

If the reserved words are specified previously in the **Authentication Method** field, this field is ignored.

|           |   |
|-----------|---|
| Data type | Unicode characters except for non-ASCII characters, but including the number sign (#), percent sign (%), and the square brackets ([ ]). |
|-----------|---|

**Token Type Local Name:** Specifies the local name of the security token type, for example, X509v3.

### Web Services Security property collection

Use this page to view a list of additional properties for the configuration.

To view this administrative console page, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators > *key\_locator\_name***. Under Additional Properties, click **Properties**.

Click **New** to create a new property.

Click **Delete** to delete a property that you specified previously.

**Name:** Specifies the name of the property.

**Value:** Specifies the value for the property.

### Web Services: Client Security Bindings collection

Use this page to view a list of client-side binding configurations for Web Services Security. These bindings are used when a Web service is a client to another Web service.

To view this administrative console page, click **Enterprise Applications > *application\_name***. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Client Security Bindings**.

**Port:** Specifies the port used to send and receive messages from a server.

**Web Service:** Specifies the name of the Web service.

**Request Sender Binding:** Specifies the binding configuration used to send request messages to the request receiver.

Click **Edit** to configure the signing information, encryption information, and the login bindings for the request sender and view a listing of key locators in the key store.

The request sender binding information specified for the client must match the request receiver binding information specified for the server.

**Response Receiver Binding:** Specifies the binding configuration used to receive response messages from response sender.

Click **Edit** to configure the signing and encryption information, and to view a list of trust anchors, intermediate certificates found in the collection certificate store, and the key locator objects for the response receiver.

The response receiver binding information specified for the client must match the response sender binding information specified for the server.

### **Request Sender Binding collection**

Use this page to specify the binding configuration to send request messages for Web Services Security.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Client Security Bindings**. Under Request Sender Binding, click **Edit**.

**Signing Information:** Specifies the configuration for the signing parameters. Signing information is used to sign and validate parts of the message including the body and time stamp.

These parameters also can be used for X.509 certificate validation when "Authentication Method" is "IDAssertion" and "ID Type" is "X509Certificate" in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

**Encryption Information:** Specifies the configuration for the encrypting and decrypting parameters. Encryption information is used for encrypting and decrypting various parts of a message including the body and user name token.

**Key Locators:** Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

**Login Mappings:** Specifies a list of configurations for validating tokens within incoming messages.

Login mappings map the authentication method to the Java Authentication and Authorization Service (JAAS) configuration.

To configure JAAS, use the administrative console and click **Security > JAAS Configuration**.

### **Login Bindings configuration settings**

Use this page to configure the encryption and decryption parameters.

The pluggable token uses the Java Authentication and Authorization Service (JAAS) CallbackHandler (javax.security.auth.callback.CallbackHandler) interface to generate the token that is inserted into the message. The following is a listing of Callback support implementations:

#### **com.ibm.wsspi.wssecurity.auth.callback.BinaryTokenCallback**

This implementation is used for generating binary tokens inserted as <wsse:BinarySecurityToken/@ValueType> in the message.

**javax.security.auth.callback.NameCallback and  
javax.security.auth.callback.NameCallback**

This implementation is used for generating user name tokens inserted as <wsse:UsernameToken> in the message.

**com.ibm.wsspi.wssecurity.auth.callback.XMLTokenSenderCallback**

This implementation is used to generate XML tokens and is inserted as the <SAML: Assertion> element in the message.

**com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback**

This implementation is used to obtain properties specified in the binding file.

To view this administrative console page, click **Applications > Enterprise Applications > application\_name**. Under Related Items, click **Web Module > URI\_file\_name > Web Services: Client Security Bindings >**. Under Request Sender Bindings, click **Edit > Login Binding**.

If the encryption information is not available, select **None**.

If the encryption information is available, select **Dedicated Login Binding** and specify the configuration in the following fields:

**Authentication Method:** Specifies the unique name for the authentication method.

**Callback Handler:** Specifies the name of the callback handler. The callback handler must implement the javax.security.auth.callback.CallbackHandler interface.

**Basic Auth User ID:** Specifies the user name for basic authentication. Basic Auth provides the capability to define a user ID and password in the binding file.

**Basic Auth Password:** Specifies the password for basic authentication.

**Token Type URI:** Specifies the URI for the token type. This information is inserted as <wsse:BinarySecurityToken>/ValueType for the XML token <SAML: Assertion>.

**Token Type Local Name:** Specifies the local name for the token type. This information is inserted as <wsse:BinarySecurityToken>/ValueType for the XML token <SAML: Assertion>.

## **Signing Information configuration settings**

Use this page to configure new signing parameters.

To view this administrative console page, click **Enterprise Applications > application\_name**. Under Related Items, click **Web Modules > URI\_file\_name > Web Services: Client Security Bindings**. In the Request Sender Binding column, click **Edit > Signing Information**.

If the signing information is not available, select **None**.

If the signing information is available, select **Dedicated Signing Information** and specify the configuration in the following fields:

**Signature Method:** Specifies the algorithm URI of the signature method. This method contains the actual value of the digital signature encoded using base64.

The following algorithms are supported:

- <http://www.w3.org/2000/09/xmldsig#rsa-sha1>
- <http://www.w3.org/2000/09/xmldsig#dsa-sha1>

**Digest Method:** Specifies the algorithm URI of the digest method.

The <http://www.w3.org/2000/09/xmldsig#sha1> algorithm is supported.

**Canonicalization Method:** Specifies the algorithm URI of the canonicalization method.

The following algorithms are supported:

- <http://www.w3.org/2001/10/xml-exc-c14n#>
- <http://www.w3.org/2001/10/xml-exc-c14n#WithComments>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments>

**Signing Key:** Specifies the key information that is used for signing. These fields are ignored in receiver-side configuration.

If the signing key is not available, select **None**.

**Certificate Path:** Specifies the settings for the certificate path validation. When you select **Trust Any**, this validation is skipped and all the incoming certificates are trusted. These fields are ignored in sender-side configuration.

If there is not a certificate path, select **None**.

If there is a certificate path, select **Trust Any** or select a Trust Anchor and a Certificate Store.

#### **Trust Anchor**

The selections available for Trust Anchor are specified by clicking **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

#### **Certificate Store**

The selections available for the Collection Store are specified by clicking **Servers > Application Servers *server\_name***. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

### **Response Receiver Binding collection**

Use this page to specify the binding configuration for receiver response messages for Web Services Security.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**. Under Response Sender Binding, click **Edit**.

**Signing Information:** Specifies the configuration for the signing parameters. Signing information is used to sign and validate parts of the message including the body and time stamp.

These parameters also can be used for X.509 certificate validation when "Authentication Method" is "IDAssertion" and "ID Type" is "X509Certificate" in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

**Encryption Information:** Specifies the configuration for the encrypting and decrypting parameters. Encryption information is used for encrypting and decrypting various parts of a message including the body and user name token.

**Trust Anchors:** Specifies a list of key store objects that contain the trusted root certificates, self-signed or issued by a certificate authority (CA).

The certificate authority authenticates a user and issues a certificate. After the certificate is issued, the key store objects, which contain these certificates, use the certificate for certificate path or certificate chain validation of incoming X.509-formatted security tokens.

**Collection Certificate Store:** Specifies a list of the untrusted, intermediate certificate files.

The collection certificate store contains a chain of untrusted, intermediate certificates. The CertPath API attempts to validate these certificates, which are based on the trust anchor.

**Key Locators:** Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

### **Web Services: Server Security Bindings collection**

Use this page to view a list of server-side binding configurations for Web Services Security.

To view this administrative console page, click **Enterprise Applications** > <application\_name>. Under Related Items, click **Web Module** > <URI\_file\_name> > **Web Services: Server Security Bindings**.

**Port:** Specifies the port in which messages are received from the request sender. .

**Web Service:** Specifies the name of the Web service.

**Request Receiver Binding:** Specifies the binding configuration used to receive request messages from the request sender.

Click **Edit** to configure the signing information and encryption information and view a listing of trust anchors, intermediate certificates in the collection certificate store, key locators, trusted ID evaluators, and login mappings.

The request receiver binding information specified for the server must match the request sender binding information specified for the client.

**Response Sender Binding:** Specifies the binding configuration used to send request messages to the response receiver.

Click **Edit** to configure the signing and encryption information, and to view a list of key locator objects for the response sender.

The response sender binding information specified for the server must match the response receiver binding information specified for the client.

### **Request Receiver Binding collection**

Use this page to specify the binding configuration to receive request messages for Web Services Security.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name* > Related Items, click Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit**.

**Signing Information:** Specifies the configuration for the signing parameters. Signing information is used to sign and validate parts of a message including the body, time stamp, and user name token.

These parameters also can be used for X.509 certificate validation when **Authentication Method** is IDAssertion and **ID Type** is X509Certificate in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

**Encryption Information:** Specifies the configuration for the encrypting and decrypting parameters. This configuration is used to encrypt and decrypt parts of the message include the body and user name token.

**Trust Anchors:** Specifies a list of key store objects that contain the trusted, root certificates issued by a certificate authority (CA).

The certificate authority authenticates a user and issues a certificate. The CertPath API uses the certificate to validate the certificate chain of incoming, X.509-formatted security tokens or trusted, self-signed certificates.

**Collection Certificate Store:** Specifies a list of the untrusted, intermediate certificate files.

The collection certificate store contains a chain of untrusted, intermediate certificates. The CertPath API attempts to validate these certificates, which are based on the trust anchor.

**Key Locators:** Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

**Trusted ID Evaluators:** Specifies a list of trusted ID evaluators that determine whether to trust the identity-asserting authority or message sender.

The trusted ID evaluators are used to authenticate additional identities from one server to another server. For example, a client sends the identity of user A to server 1 for authentication. Server 1 calls downstream to server 2, asserts the identity of user A, and includes the user ID and password of server 1. Server 2 attempts to establish trust with server 1 by authenticating its user ID and password and checking the trust based on the TrustedIDEvaluator implementation. If the authentication process and the trust check are successful, server 2 trusts that server 1 authenticated user A and a credential is created for user A on server 2 to invoke the request.

**Login Mappings:** Specifies a list of configurations for validating tokens within incoming messages.

Login mappings map the authentication method to the Java Authentication and Authorization Service (JAAS) configuration.

To configure JAAS, use the administrative console and click **Security > JAAS Configuration**.

### Signing Information collection

Use this page to view a list of signing parameters. Signing information is used to sign and validate parts of a message including the body, time stamp, and user name token. These parameters also can be used for X.509 certificate validation when **Authentication Method** is IDAssertion and **ID Type** is X509Certificate in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

To view this administrative console page, click **Applications > Enterprise Applications > application\_name**. Under Related Items, click **Web Modules > URI\_file\_name > Web Services: Server Security Bindings**. In the Request Receiver Binding column, click **Edit > Signing Information**.

Click **New** to create a signing parameter.

Click **Delete** to delete a signing parameter.

**Signature Method:** Specifies the unique name of the signature method.

### Signing Information configuration settings

Use this page to configure new signing parameters.

The specifications listed on this page for the signature method, digest method, and canonicalization method are located in the W3C document entitled, "XMLSignature Syntax and Specification: W3C Recommendation 12 Feb 2002".

To view this administrative console page, click **Enterprise Applications > application\_name**. Under Related Items, click **Web Modules > URI\_file\_name > Web Services: Server Security Bindings**. In the Request Receiver Binding column, click **Edit > Signing Information > New**.

Click **New** to create a signing parameter.

Click **Delete** to delete a signing parameter.

**Signature Method:** Specifies the algorithm URI of the signature method. This method contains the actual value of the digital signature encoded using base64.

The following algorithms are supported:

- <http://www.w3.org/2000/09/xmlsig#rsa-sha1>
- <http://www.w3.org/2000/09/xmlsig#dsa-sha1>

**Digest Method:** Specifies the algorithm URI of the digest method.

The <http://www.w3.org/2000/09/xmlsig#sha1> algorithm is supported.

**Canonicalization Method:** Specifies the algorithm URI of the canonicalization method.

The following algorithms are supported:

- <http://www.w3.org/2001/10/xml-exc-c14n#>
- <http://www.w3.org/2001/10/xml-exc-c14n#WithComments>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments>

**Signing Key:** Specifies the key information that is used for signing. These fields are ignored in receiver-side configuration.

If you specify a **Key Name** and a **Key Locator Reference**, select **None** for the Certificate Path.

**Certificate Path:** Specifies the settings for the certificate path validation. When you select **Trust Any**, this validation is skipped and all the incoming certificates are trusted. These fields are ignored in sender-side configuration.

If you click **Trust Any** or select a **Trust Anchor** and a **Certificate Store**, select **None** for the Signing Key in the previous field.

#### **Trust Anchor**

The selections available for Trust Anchor are specified by clicking **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

#### **Certificate Store**

The selections available for the Collection Store are specified by clicking **Servers > Application Servers > *server\_name***. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

### **Encryption Information collection**

Use this page to specify the configuration for the encrypting and decrypting parameters. This configuration is used to encrypt and decrypt parts of the message including the body and user name token.

To view this administrative console page, click **Enterprise Applications > *application\_name***. Under Related Items, Click **Web Module**. Under Additional Properties, click **Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit > Encryption Information**.

Click **New** to create an encryption method.

Click **Delete** to delete an encryption method.

**Encryption Information:** Specifies the name of the encryption information.

### **Encryption information configuration settings**

Use this page to configure the encryption and decryption parameters. You can use these parameters to encrypt and decrypt various parts of the message including the body and user name token.



To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit > Encryption Information > New**.

**Encryption Information Name:** Specifies the name for the encryption information.

**Key Locator Reference:** Specifies the name used to reference the key locator.

To specify key locator references, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

**Encryption Key Name:** Specifies the name of the encryption key, which is resolved to the actual key by the specified key locator.

**Key Encryption Algorithm:** Specifies the algorithm URI of the key encryption method.

The following algorithms are supported:

- [http://www.w3.org/2001/04/xmlenc#rsa-1\\_5](http://www.w3.org/2001/04/xmlenc#rsa-1_5)
- <http://www.w3.org/2001/04/xmlenc#kw-tripledes>

**Data Encryption Algorithm:** Specifies the algorithm URI of the data encryption method.

The <http://www.w3.org/2001/04/xmlenc#tripledes-cbc> algorithm is supported.

## Response Sender Binding collection

Use this page to specify the binding configuration for sender response messages for Web Services Security.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**. Under Response Sender Binding, click **Edit**.

**Signing Information:** Specifies the configuration for the signing parameters.

These parameters also can be used for X.509 certificate validation when "Authentication Method" is "IDAssertion" and "ID Type" is "X509Certificate" in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

**Encryption Information:** Specifies the configuration for the encrypting and decrypting parameters.

**Key Locators:** Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

## Encryption information configuration settings

Use this page to configure the encryption and decryption parameters.

The specifications listed on this page for the signature method, digest method, and canonicalization method are located in the W3C document entitled, "XML Encryption Syntax and Processing: W3C Recommendation 10 Dec 2002".

To view this administrative console page, click **Applications > Enterprise Applications > application\_name**. Under Related Items, click **Web Module > URI\_file\_name > Web Services: Server Security Bindings**. Under Response Sender Binding, click **Edit > Encryption Information**.

If the encryption information is not available, select **None**.

If the encryption information is available, select **Dedicated Encryption Information** and specify the configuration in the following fields:

**Encryption Information Name:** Specifies the name for the encryption information.

**Key Locator Reference:** Specifies the name used to reference the key locator.

To specify key locator references, click **Servers > Application Servers > server\_name**. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

**Encryption Key Name:** Specifies the name of the encryption key, which is resolved to the actual key by the specified key locator.

**Key Encryption Algorithm:** Specifies the algorithm URI of the key encryption method.

The following algorithms are supported:

- [http://www.w3.org/2001/04/xmlenc#rsa-1\\_5](http://www.w3.org/2001/04/xmlenc#rsa-1_5)
- <http://www.w3.org/2001/04/xmlenc#kw-tripledes>

**Data Encryption Algorithm:** Specifies the algorithm URI of the data encryption method.

The <http://www.w3.org/2001/04/xmlenc#tripledes-cbc> algorithm is supported.

## View Web Services Server Deployment Descriptor

Use this page view your server deployment descriptor settings.

To view this administrative console page, click **Applications > Enterprise Applications > application\_name**. Under Related Items, click **Web Modules > URI\_file\_name > View Web Services Server Deployment Descriptor**.

WebSphere Application Server, Version 5.02 has two levels of bindings: application-level and server-level. WebSphere Application Server Network Deployment, Version 5.02 has three levels of bindings: application-level, server-level, and cell-level. The information in the following implementation descriptions indicate how to configure your application-level bindings. To configure the server-level bindings, which are the defaults, click **Servers > Application Servers > server\_name**. Under Related Items, click **Web Services: Default bindings for Web Services Security**. To configure the cell-level bindings, click **Security > Web Services**.

- Request digital signature verification
- Request decryption

- BasicAuth authentication
- Identity (ID) Assertion authentication with the BasicAuth TrustMode
- Identity (ID) Assertion authentication with the Signature TrustMode
- Response signing
- Response encryption

### Request digital signature verification

If the integrity constraints (signature required) are defined, verify that you have configured the signing information in the binding files.

To configure the signing parameters, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > <URI\_file\_name> Web Services: Server Security Bindings**. In the Request Receiver Binding column, click **Edit > Signing Information**.

To configure the trust anchor, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

To configure the collection certificate store, click **Servers > Application Servers > *server\_name***. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

To configure the key locators, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

### Request decryption

If the confidentiality constraints (encryption) are specified, verify that the encryption information is defined.

To configure the encryption information parameters, click **Enterprise Applications > *application\_name***. Under Related Items, Click **Web Module**. Under Additional Properties, click **Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit > Encryption Information**.

To configure the key locators, click click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

### BasicAuth authentication

If BasicAuth is configured as the required security token, specify the CallbackHandler in the binding file to collect the basic authentication data. The following is a listing of Callback support implementations:

#### **com.ibm.wsspi.wssecurity.auth.callback.GuiPromptCallbackHandler**

The implementation prompts for BasicAuth information (user name and password) in a GUI panel.

#### **com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler**

This implementation reads the BasicAuth information from the binding file.

### **com.ibm.wsspi.wssecurity.auth.callback.StdPromptCallbackHandler**

This implementation prompts for a user name and password using standard in (stdin).

To configure the login mapping information, click **Server > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings**.

### **Identity (ID) Assertion authentication with the BasicAuth TrustMode**

Configure a login binding in the bindings file with a `com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler` implementation. Specify a BasicAuth user ID and password that a TrustedIDEvaluator on a downstream server will trust.

To configure the login mapping information, click **Server > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings**.

### **Identity (ID) Assertion authentication with the Signature TrustMode**

Configure the signing information in the bindings file with a signing key pointing to a key locator. The key locator contains the X.509 certificate that will be trusted by the downstream server.

To configure the login mapping information, click **Server > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings**.

WSLogin is the name of the login configuration used by the Java Authentication and Authorization Service (JAAS). To configure JAAS, click **Security > JAAS Configuration > Application Logins**.

The value of the `<TrustedIDEvaluatorRef>` tag in the binding must match the value of the `<TrustedIDEvaluator>` name.

To configure the trusted ID evaluators, click **Servers > Application Servers > *server\_name***. Under Additional Services, click **Web Services: Default bindings for Web Services Security > Trusted ID Evaluators**.

### **Response signing**

If the integrity constraints (digital signature) are defined, verify that you have the signing information configured in the binding files.

To specify the signing information, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* Web Services: Server Security Bindings**. In the Request Receiver Binding column, click **Edit > Signing Information**.

To configure the key locators, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

### **Response encryption**

If the confidentiality constraints (encryption) are specified, verify that the encryption information is defined.

To specify the encryption information, click **Enterprise Applications** > *application\_name*. Under Related Items, Click **Web Module**. Under Additional Properties, click **Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit** > **Encryption Information**.

To configure the key locators, click **Servers** > **Application Servers** > *server\_name*. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Key Locators**.

## View Web Services Client Deployment Descriptor

Use this page view your client deployment descriptor.

To view this administrative console page, click **Applications** > **Enterprise Applications** > *application\_name*. Under Related Items, click **Web Modules** > *URI\_file\_name* > **View Web Services Client Deployment Descriptor**.

WebSphere Application Server, Version 5.02 has two levels of bindings: application-level and server-level. WebSphere Application Server Network Deployment, Version 5.02 has three levels of bindings: application-level, server-level, and cell-level. The information in the following implementation descriptions indicate how to configure your application-level bindings. To configure the server-level bindings, which are the defaults, click **Servers** > **Application Servers** > *server\_name*. Under Related Items, click **Web Services: Default bindings for Web Services Security**. To configure the cell-level bindings, click **Security** > **Web Services**.

**Note:** If the Web server is acting as a client, the default bindings are used.

If you are using any of the following configurations, verify that the deployment descriptor is configured properly:

- Request signing
- Request encryption
- BasicAuth authentication
- Identity (ID) Assertion authentication with the BasicAuth TrustMode
- Identity (ID) Assertion authentication with the Signature TrustMode
- Response digital signature verification
- Response decryption

### Request signing

If the Integrity constraints (digital signature) are specified, verify that you have configured the signing information in the binding files.

To configure the signing parameters, click **Applications** > **Enterprise Applications** > *application\_name*. Under Related Items, click **Web Modules** > *URI\_file\_name* **Web Services: Client Security Bindings**. In the Response Receiver Binding column, click **Edit** > **Signing Information** > **New**.

To configure the key locators, click **Servers** > **Application Servers** > *server\_name*. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Key Locators**.

## Request encryption

If the confidentiality constraints (encryption) are specified, verify that you have configured the encryption information in the binding files.

To configure the encryption parameters, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* Web Services: Client Security Bindings**. In the Response Receiver Binding column, click **Edit > Encryption Information > New**.

To configure the key locators, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

## BasicAuth authentication

If BasicAuth is configured as the required security token, specify the CallbackHandler in the binding file to collect the basic authentication data. The following is a listing of Callback support implementations:

### **com.ibm.wsspi.wssecurity.auth.callback.GuiPromptCallbackHandler**

The implementation prompts for BasicAuth information (user name and password) in a GUI panel.

### **com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler**

This implementation reads the BasicAuth information from the binding file.

### **com.ibm.wsspi.wssecurity.auth.callback.StdPromptCallbackHandler**

This implementation prompts for a user name and password using standard in (stdin).

To configure the login binding information, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Module > *URI\_file\_name* Web Services: Client Security Bindings**. Under Request Sender Bindings, click **Edit > Login Binding**.

## Identity (ID) Assertion authentication with BasicAuth TrustMode

Configure a login binding in the bindings file with a `com.ibm.wsspi.wssecurity.auth.callback.NonPromptCallbackHandler` implementation. Specify a BasicAuth user ID and password that a TrustedIDEvaluator on a downstream server will trust.

To configure the login binding information, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Module > *URI\_file\_name* Web Services: Client Security Bindings**. Under Request Sender Bindings, click **Edit > Login Binding**.

## Identity (ID) Assertion authentication with the Signature TrustMode

Configure the signing information in the bindings file with a signing key pointing to a key locator. The key locator contains the X.509 certificate that will be trusted by the downstream server.

To configure ID assertion, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Login Mappings > IDAssertion**.

To configure the login binding information, click **Applications > Enterprise Applications > application\_name**. Under Related Items, click **Web Module > URI\_file\_name Web Services: Client Security Bindings**. Under Request Sender Bindings, click **Edit > Login Binding**.

### **Response digital signature verification**

If the integrity constraints (signature required) are defined, verify that you have configured the signing information in the binding files.

To configure the signing parameters, click **Applications > Enterprise Applications > application\_name**. Under Related Items, click **Web Modules > URI\_file\_name > Web Services: Client Security Bindings**. In the Response Receiver Binding column, click **Edit > Signing Information > New**.

To configure the trust anchors, click **Servers > Application Servers > server\_name**. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors > New**.

To configure the collection certificate store, click **Servers > Application Servers > server\_name**. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store > New**.

### **Response decryption**

If the confidentiality constraints (encryption) are specified, verify that the encryption information is defined.

To configure the encryption information, click **Applications > Enterprise Applications > application\_name**. Under Related Items, click **Web Modules > <URI\_file\_name> Web Services: Client Security Bindings**. In the Response Receiver Binding column, click **Edit > Encryption Information > New**.

To configure the key locators, click **Servers > Application Servers > server\_name**. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

## **Web Services: Server Security Bindings collection**

Use this page to view a list of server-side binding configurations for Web Services Security.

To view this administrative console page, click **Enterprise Applications > <application\_name>**. Under Related Items, click **Web Module > <URI\_file\_name> > Web Services: Server Security Bindings**.

### **Port**

Specifies the port in which messages are received from the request sender. .

### **Web Service**

Specifies the name of the Web service.

### **Request Receiver Binding**

Specifies the binding configuration used to receive request messages from the request sender.

Click **Edit** to configure the signing information and encryption information and view a listing of trust anchors, intermediate certificates in the collection certificate store, key locators, trusted ID evaluators, and login mappings.

The request receiver binding information specified for the server must match the request sender binding information specified for the client.

### **Response Sender Binding**

Specifies the binding configuration used to send request messages to the response receiver.

Click **Edit** to configure the signing and encryption information, and to view a list of key locator objects for the response sender.

The response sender binding information specified for the server must match the response receiver binding information specified for the client.

### **Request Receiver Binding collection**

Use this page to specify the binding configuration to receive request messages for Web Services Security.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit**.

**Signing Information:** Specifies the configuration for the signing parameters. Signing information is used to sign and validate parts of a message including the body, time stamp, and user name token.

These parameters also can be used for X.509 certificate validation when **Authentication Method** is IDAssertion and **ID Type** is X509Certificate in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

**Encryption Information:** Specifies the configuration for the encrypting and decrypting parameters. This configuration is used to encrypt and decrypt parts of the message include the body and user name token.

**Trust Anchors:** Specifies a list of key store objects that contain the trusted, root certificates issued by a certificate authority (CA).

The certificate authority authenticates a user and issues a certificate. The CertPath API uses the certificate to validate the certificate chain of incoming, X.509-formatted security tokens or trusted, self-signed certificates.

**Collection Certificate Store:** Specifies a list of the untrusted, intermediate certificate files.

The collection certificate store contains a chain of untrusted, intermediate certificates. The CertPath API attempts to validate these certificates, which are based on the trust anchor.

**Key Locators:** Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key



locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

**Trusted ID Evaluators:** Specifies a list of trusted ID evaluators that determine whether to trust the identity-asserting authority or message sender.

The trusted ID evaluators are used to authenticate additional identities from one server to another server. For example, a client sends the identity of user A to server 1 for authentication. Server 1 calls downstream to server 2, asserts the identity of user A, and includes the user ID and password of server 1. Server 2 attempts to establish trust with server 1 by authenticating its user ID and password and checking the trust based on the TrustedIDEvaluator implementation. If the authentication process and the trust check are successful, server 2 trusts that server 1 authenticated user A and a credential is created for user A on server 2 to invoke the request.

**Login Mappings:** Specifies a list of configurations for validating tokens within incoming messages.

Login mappings map the authentication method to the Java Authentication and Authorization Service (JAAS) configuration.

To configure JAAS, use the administrative console and click **Security > JAAS Configuration**.

### **Signing Information collection**

Use this page to view a list of signing parameters. Signing information is used to sign and validate parts of a message including the body, time stamp, and user name token. These parameters also can be used for X.509 certificate validation when **Authentication Method** is IDAssertion and **ID Type** is X509Certificate in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**. In the Request Receiver Binding column, click **Edit > Signing Information**.

Click **New** to create a signing parameter.

Click **Delete** to delete a signing parameter.

**Signature Method:** Specifies the unique name of the signature method.

### **Signing Information configuration settings**

Use this page to configure new signing parameters.

The specifications listed on this page for the signature method, digest method, and canonicalization method are located in the W3C document entitled, "XMLSignature Syntax and Specification: W3C Recommendation 12 Feb 2002".

To view this administrative console page, click **Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**. In the Request Receiver Binding column, click **Edit > Signing Information > New**.

Click **New** to create a signing parameter.

Click **Delete** to delete a signing parameter.

**Signature Method:** Specifies the algorithm URI of the signature method. This method contains the actual value of the digital signature encoded using base64.

The following algorithms are supported:

- <http://www.w3.org/2000/09/xmlsig#rsa-sha1>
- <http://www.w3.org/2000/09/xmlsig#dsa-sha1>

**Digest Method:** Specifies the algorithm URI of the digest method.

The <http://www.w3.org/2000/09/xmlsig#sha1> algorithm is supported.

**Canonicalization Method:** Specifies the algorithm URI of the canonicalization method.

The following algorithms are supported:

- <http://www.w3.org/2001/10/xml-exc-c14n#>
- <http://www.w3.org/2001/10/xml-exc-c14n#WithComments>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments>

**Signing Key:** Specifies the key information that is used for signing. These fields are ignored in receiver-side configuration.

If you specify a **Key Name** and a **Key Locator Reference**, select **None** for the Certificate Path.

**Certificate Path:** Specifies the settings for the certificate path validation. When you select **Trust Any**, this validation is skipped and all the incoming certificates are trusted. These fields are ignored in sender-side configuration.

If you click **Trust Any** or select a **Trust Anchor** and a **Certificate Store**, select **None** for the Signing Key in the previous field.

#### **Trust Anchor**

The selections available for Trust Anchor are specified by clicking **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

#### **Certificate Store**

The selections available for the Collection Store are specified by clicking **Servers > Application Servers > *server\_name***. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

### **Encryption Information collection**

Use this page to specify the configuration for the encrypting and decrypting parameters. This configuration is used to encrypt and decrypt parts of the message including the body and user name token.

To view this administrative console page, click **Enterprise Applications** > *application\_name*. Under Related Items, Click **Web Module**. Under Additional Properties, click **Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit** > **Encryption Information**.

Click **New** to create an encryption method.

Click **Delete** to delete an encryption method.

**Encryption Information:** Specifies the name of the encryption information.

### **Encryption information configuration settings**

Use this page to configure the encryption and decryption parameters. You can use these parameters to encrypt and decrypt various parts of the message including the body and user name token.

To view this administrative console page, click **Applications** > **Enterprise Applications** > *application\_name*. Under Related Items, click **Web Module** > *URI\_file\_name* > **Web Services: Server Security Bindings**. Under Request Receiver Binding, click **Edit** > **Encryption Information** > **New**.

**Encryption Information Name:** Specifies the name for the encryption information.

**Key Locator Reference:** Specifies the name used to reference the key locator.

To specify key locator references, click **Servers** > **Application Servers** > *server\_name*. Under Additional Properties, click **Web Services: Default bindings for Web Services Security** > **Key Locators**.

**Encryption Key Name:** Specifies the name of the encryption key, which is resolved to the actual key by the specified key locator.

**Key Encryption Algorithm:** Specifies the algorithm URI of the key encryption method.

The following algorithms are supported:

- [http://www.w3.org/2001/04/xmlenc#rsa-1\\_5](http://www.w3.org/2001/04/xmlenc#rsa-1_5)
- <http://www.w3.org/2001/04/xmlenc#kw-tripledes>

**Data Encryption Algorithm:** Specifies the algorithm URI of the data encryption method.

The <http://www.w3.org/2001/04/xmlenc#tripledes-cbc> algorithm is supported.

### **Response Sender Binding collection**

Use this page to specify the binding configuration for sender response messages for Web Services Security.

To view this administrative console page, click **Applications** > **Enterprise Applications** > *application\_name*. Under Related Items, click **Web Modules** > *URI\_file\_name* > **Web Services: Server Security Bindings**. Under Response Sender Binding, click **Edit**.

**Signing Information:** Specifies the configuration for the signing parameters.

These parameters also can be used for X.509 certificate validation when "Authentication Method" is "IDAssertion" and "ID Type" is "X509Certificate" in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

**Encryption Information:** Specifies the configuration for the encrypting and decrypting parameters.

**Key Locators:** Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

### **Encryption information configuration settings**

Use this page to configure the encryption and decryption parameters.

The specifications listed on this page for the signature method, digest method, and canonicalization method are located in the W3C document entitled, "XML Encryption Syntax and Processing: W3C Recommendation 10 Dec 2002".

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Server Security Bindings**. Under Response Sender Binding, click **Edit > Encryption Information**.

If the encryption information is not available, select **None**.

If the encryption information is available, select **Dedicated Encryption Information** and specify the configuration in the following fields:

**Encryption Information Name:** Specifies the name for the encryption information.

**Key Locator Reference:** Specifies the name used to reference the key locator.

To specify key locator references, click **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Key Locators**.

**Encryption Key Name:** Specifies the name of the encryption key, which is resolved to the actual key by the specified key locator.

**Key Encryption Algorithm:** Specifies the algorithm URI of the key encryption method.

The following algorithms are supported:

- [http://www.w3.org/2001/04/xmlenc#rsa-1\\_5](http://www.w3.org/2001/04/xmlenc#rsa-1_5)
- <http://www.w3.org/2001/04/xmlenc#kw-tripledes>

**Data Encryption Algorithm:** Specifies the algorithm URI of the data encryption method.

The <http://www.w3.org/2001/04/xmlenc#tripledes-cbc> algorithm is supported.

## Web Services: Client Security Bindings collection

Use this page to view a list of client-side binding configurations for Web Services Security. These bindings are used when a Web service is a client to another Web service.

To view this administrative console page, click **Enterprise Applications** > *application\_name*. Under Related Items, click **Web Module** > *URI\_file\_name* > **Web Services: Client Security Bindings**.

### Port

Specifies the port used to send and receive messages from a server.

### Web Service

Specifies the name of the Web service.

### Request Sender Binding

Specifies the binding configuration used to send request messages to the request receiver.

Click **Edit** to configure the signing information, encryption information, and the login bindings for the request sender and view a listing of key locators in the key store.

The request sender binding information specified for the client must match the request receiver binding information specified for the server.

### Response Receiver Binding

Specifies the binding configuration used to receive response messages from response sender.

Click **Edit** to configure the signing and encryption information, and to view a list of trust anchors, intermediate certificates found in the collection certificate store, and the key locator objects for the response receiver.

The response receiver binding information specified for the client must match the response sender binding information specified for the server.

### Request Sender Binding collection

Use this page to specify the binding configuration to send request messages for Web Services Security.

To view this administrative console page, click **Applications** > **Enterprise Applications** > *application\_name*. Under Related Items, click **Web Modules** > *URI\_file\_name* > **Web Services: Client Security Bindings**. Under Request Sender Binding, click **Edit**.

**Signing Information:** Specifies the configuration for the signing parameters. Signing information is used to sign and validate parts of the message including the body and time stamp.

These parameters also can be used for X.509 certificate validation when "Authentication Method" is "IDAssertion" and "ID Type" is "X509Certificate" in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

**Encryption Information:** Specifies the configuration for the encrypting and decrypting parameters. Encryption information is used for encrypting and decrypting various parts of a message including the body and user name token.

**Key Locators:** Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

**Login Mappings:** Specifies a list of configurations for validating tokens within incoming messages.

Login mappings map the authentication method to the Java Authentication and Authorization Service (JAAS) configuration.

To configure JAAS, use the administrative console and click **Security > JAAS Configuration**.

### **Login Bindings configuration settings**

Use this page to configure the encryption and decryption parameters.

The pluggable token uses the Java Authentication and Authorization Service (JAAS) CallbackHandler (javax.security.auth.callback.CallbackHandler) interface to generate the token that is inserted into the message. The following is a listing of Callback support implementations:

**com.ibm.wsspi.wssecurity.auth.callback.BinaryTokenCallback**

This implementation is used for generating binary tokens inserted as `<wsse:BinarySecurityToken/@ValueType>` in the message.

**javax.security.auth.callback.NameCallback and  
javax.security.auth.callback.NameCallback**

This implementation is used for generating user name tokens inserted as `<wsse:UsernameToken>` in the message.

**com.ibm.wsspi.wssecurity.auth.callback.XMLTokenSenderCallback**

This implementation is used to generate XML tokens and is inserted as the `<SAML: Assertion>` element in the message.

**com.ibm.wsspi.wssecurity.auth.callback.PropertyCallback**

This implementation is used to obtain properties specified in the binding file.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Module > *URI\_file\_name* > Web Services: Client Security Bindings >**. Under Request Sender Bindings, click **Edit > Login Binding**.

If the encryption information is not available, select **None**.

If the encryption information is available, select **Dedicated Login Binding** and specify the configuration in the following fields:

**Authentication Method:** Specifies the unique name for the authentication method.

**Callback Handler:** Specifies the name of the callback handler. The callback handler must implement the javax.security.auth.callback.CallbackHandler interface.

**Basic Auth User ID:** Specifies the user name for basic authentication. Basic Auth provides the capability to define a user ID and password in the binding file.

**Basic Auth Password:** Specifies the password for basic authentication.

**Token Type URI:** Specifies the URI for the token type. This information is inserted as <wsse:BinarySecurityToken>/ValueType for the XML token <SAML: Assertion>.

**Token Type Local Name:** Specifies the local name for the token type. This information is inserted as <wsse:BinarySecurityToken>/ValueType for the XML token <SAML: Assertion>.

## Signing Information configuration settings

Use this page to configure new signing parameters.

To view this administrative console page, click **Enterprise Applications > application\_name**. Under Related Items, click **Web Modules > URI\_file\_name > Web Services: Client Security Bindings**. In the Request Sender Binding column, click **Edit > Signing Information**.

If the signing information is not available, select **None**.

If the signing information is available, select **Dedicated Signing Information** and specify the configuration in the following fields:

**Signature Method:** Specifies the algorithm URI of the signature method. This method contains the actual value of the digital signature encoded using base64.

The following algorithms are supported:

- <http://www.w3.org/2000/09/xmlsig#rsa-sha1>
- <http://www.w3.org/2000/09/xmlsig#dsa-sha1>

**Digest Method:** Specifies the algorithm URI of the digest method.

The <http://www.w3.org/2000/09/xmlsig#sha1> algorithm is supported.

**Canonicalization Method:** Specifies the algorithm URI of the canonicalization method.

The following algorithms are supported:

- <http://www.w3.org/2001/10/xml-exc-c14n#>
- <http://www.w3.org/2001/10/xml-exc-c14n#WithComments>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315>
- <http://www.w3.org/TR/2001/REC-xml-c14n-20010315#WithComments>

**Signing Key:** Specifies the key information that is used for signing. These fields are ignored in receiver-side configuration.

If the signing key is not available, select **None**.

**Certificate Path:** Specifies the settings for the certificate path validation. When you select **Trust Any**, this validation is skipped and all the incoming certificates are trusted. These fields are ignored in sender-side configuration.

If there is not a certificate path, select **None**.

If there is a certificate path, select **Trust Any** or select a Trust Anchor and a Certificate Store.

### **Trust Anchor**

The selections available for Trust Anchor are specified by clicking **Servers > Application Servers > *server\_name***. Under Additional Properties, click **Web Services: Default bindings for Web Services Security > Trust Anchors**.

### **Certificate Store**

The selections available for the Collection Store are specified by clicking **Servers > Application Servers *server\_name***. Under Related Items, click **Web Services: Default bindings for Web Services Security > Collection Certificate Store**.

### **Response Receiver Binding collection**

Use this page to specify the binding configuration for receiver response messages for Web Services Security.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***. Under Related Items, click **Web Modules > *URI\_file\_name* > Web Services: Server Security Bindings**. Under Response Sender Binding, click **Edit**.

**Signing Information:** Specifies the configuration for the signing parameters. Signing information is used to sign and validate parts of the message including the body and time stamp.

These parameters also can be used for X.509 certificate validation when "Authentication Method" is "IDAssertion" and "ID Type" is "X509Certificate" in the server-level configuration. In such cases, you must fill-in the "Certificate Path" fields only.

**Encryption Information:** Specifies the configuration for the encrypting and decrypting parameters. Encryption information is used for encrypting and decrypting various parts of a message including the body and user name token.

**Trust Anchors:** Specifies a list of key store objects that contain the trusted root certificates, self-signed or issued by a certificate authority (CA).

The certificate authority authenticates a user and issues a certificate. After the certificate is issued, the key store objects, which contain these certificates, use the certificate for certificate path or certificate chain validation of incoming X.509-formatted security tokens.

**Collection Certificate Store:** Specifies a list of the untrusted, intermediate certificate files.

The collection certificate store contains a chain of untrusted, intermediate certificates. The CertPath API attempts to validate these certificates, which are based on the trust anchor.

**Key Locators:** Specifies a list of key locator objects that retrieve the keys for digital signature and encryption from a key store file or a repository. The key



locator maps a name or logical name to an alias or maps an authenticated identity to a key. This is the logical name used to locate a key in a key locator implementation.

---

## Tuning Web services applications

### Before you begin

Since a Web service application is a regular J2EE application, the performance considerations are primarily the same as they are for any J2EE application. See Tuning performance for more information about analyzing and tuning J2EE applications.

You can use the Performance Monitoring Infrastructure (PMI) to measure the time required to process Web services requests. To monitor Web services application performance:

### Steps for this task

1. Enable PMI services in application server through the administrative console. Select the Web Service module, named `webServicesModule`, in step 7.
2. Monitor performance with Tivoli Performance Viewer  
In the left-hand pane of the performance view, expand the host and server and select **Web Services**. Run the Web services client application.

### Results

Measurements are available for the following items:

- Number of Web services loaded by the application server
- Number of requests received
- Number of requests dispatched to an implementation bean
- Number of requests dispatched with successful replies
- Average time in milliseconds between receiving the request and returning the reply
- Average time in milliseconds between receiving the request and dispatching it to the bean
- Average time in milliseconds between dispatch and receipt of reply from the bean
- Average time in milliseconds between receipt of reply from bean to return of result to client
- Average size of request and reply
- Average size of request
- Average size of reply

---

## Troubleshooting Web services

To work through problems with Web services, you can:

### Steps for this task

1. Trace Web services messages.

## Tracing Web services messages

### Before you begin

You can trace the XML messages exchanged between a client and the server using the **TCPMon** command tool. The **TCPMon** command redirects messages from one port to another and records them. The WebSphere Application Server listens on port 9080. To trace messages sent to the application server, the **TCPMon** command is configured to listen on port 9088 and redirect them to 9080. The client is redirected to use port 9088 to access the Web service. You should confirm that the server providing the Web service is running. The following task is performed on the machine providing the Web service.

To trace error messages in Web services:

### Steps for this task

1. Set the CLASSPATH to include the *install\_root/lib/webservices.jar* file.
2. Run the **java com.ibm.ws.webservices.engine.utils.tcpmon** command.  
A window labeled TCPMonitor displays.
3. Configure the TCPMonitor to listen on port 9088 and forward messages to port 9080.
  - a. In the **Listen Port #** field, enter 9088.
  - b. Click **Listener**
  - c. In the **TargetHostname** field, enter localhost.
  - d. In the **Target Port #** field, enter 9080.
  - e. Click **Add**.
  - f. Click on the **Port 9088** tab that displays on the top of the page.

### Results

The messages exchanged between the client and server appear in the TCPMonitor Request and Response pane.

### What to do next

Save the message data and analyze it.

## Frequently asked questions about Web services for J2EE

- What IBM development tools work with Web services for J2EE?
- Is Web services for J2EE part of the J2EE specification?
- What is the relationship between Web services for J2EE and the Web Service Invocation Framework (WSIF)?
- What is the relationship between Apache Simple Object Access Protocol (SOAP) 2.3 and Web services for J2EE?
- What is the relationship between the Apache Axis component of the Web services technology preview available with WebSphere Application Server 5.0 and Web services for J2EE?
- What standards does the Web services for J2EE component of WebSphere Application Server 5.0 support?
- Does Web services for J2EE interoperate with other SOAP implementations, like .NET?

- Why can I not use a JavaBean to implement a SOAP Java Messaging Service (JMS) service?
- Does the SOAP JMS support interoperate with other vendors?
- How does two-way messaging with SOAP JMS work? Can it support multiple clients making simultaneous requests?
- When running the WSDLJava command tool on a Web Services Description Language (WSDL) document that contains a JMS-style endpoint URL, for example, `jms:/...`, the following error displays: `java.io.IOException: WSWS3099E: Error: Emitter failer. Invalid endpoint address in port <x> in service <y>: <jms-url-string>`. How can the problem be resolved?
- The malformed URL exception displays when running a client that uses a JMS endpoint URL. How can the problem be resolved?

### **What IBM development tools work with Web services for J2EE?**

WebSphere Studio Application Developer Version 5.1 and the Assembly Toolkit Version 5.1 both support the use of Web services for J2EE. The Application Assembly Tool, included with Websphere Application Server and Websphere Studio Application Developer Versions 5.1 and earlier, do not support Web services for J2EE.

### **Is Web services for J2EE part of the J2EE specification?**

For WebSphere Application Server 5.0.2, the Web services for J2EE Version 1.0 specification is an addition to J2EE 1.3. J2EE 1.4 requires support for Web services for J2EE Version 1.1. There are minor differences between the J2EE 1.3 Version (JSR-109 Version 1.0) and the J2EE 1.4 Version (JSR-109 Version 1.1).

### **What is the relationship between Web services for J2EE and the Web Service Invocation Framework (WSIF)?**

Web services for J2EE and WSIF represent two different programming models for accessing Web services. Web services for J2EE is standard, Java-centric, and more statically bound to WSDL documents due to the use of generated stubs. WSIF directly models Web Services Description Language (WSDL) documents. WSIF is more suitable when dynamically interpreting WSDL. Future versions of WebSphere Application server will leverage both technologies to achieve dynamic, high performing standards-based Web services implementations.

### **What is the relationship between Apache Simple Object Access Protocol (SOAP) 2.3 and Web services for J2EE?**

Apache SOAP shipped with WebSphere Application Server Versions 4.0 and 5.0. It continues to co-exist with Web services for J2EE. Apache SOAP is a proprietary API and applications written for it are not portable to other SOAP implementations. Applications written for Web services for J2EE should be portable to any vendor's implementation that supports Web services for J2EE.

### **What is the relationship between the Apache Axis component of the Web services technology preview available with WebSphere Application Server 5.0 and Web services for J2EE?**

The Web services technology preview leveraged the work that IBM contributed to the Apache Axis code base. The Web services for J2EE support included with WebSphere Application Server 5.0.2 is derived from Apache Axis, but has diverged

and contains many IBM-specific features to enhance performance, scalability, reliability, interoperability, and integration with the WebSphere Application Server.

### **What standards does the Web services for J2EE component of WebSphere Application Server 5.0. support?**

The following standards are supported by the Web services for J2EE component of WebSphere Application Server 5.0:

- SOAP Version 1.1
- Web Services Description Language (WSDL) Version 1.1
- Web services for J2EE (JSR-109) Version 1.0
- Java API for XML-Based RPC (JAX-RPC) Version 1.0
- SOAP with attachments API for Java (SAAJ) Version 1.1

### **Does Web services for J2EE interoperate with other SOAP implementations, like .NET?**

Web services for J2EE intends to conform to the WS-I Basic Profile 1.0, and should interoperate with any other vendor conforming to this specification. At the time of the writing, the Basic Profile 1.0 had not been completed, so it is possible that minor incompatibilities exist.

### **Why can I not use a JavaBean to implement a SOAP Java Messaging Service (JMS) service?**

The SOAP JMS support uses Message Driven Beans (MDB) to implement the JMS endpoint. MDBs can only be used in the EJB container and delegate to an enterprise bean. If you want to use a Java bean instead of an enterprise bean to implement the service endpoint, you must create a "facade" enterprise bean that delegates to the Java bean.

### **Does the SOAP JMS support interoperate with other vendors?**

No. There is currently no specification for SOAP JMS, therefore each vendor chooses its own implementation technique.

### **How does two-way messaging with SOAP JMS work? Can it support multiple clients making simultaneous requests?**

Before a client issues a two-way request, it creates a temporary JMS queue to receive the response. This temporary queue is specified as the **replyTo** destination in the outgoing JMS request message. After the server processes the request, it directs the response to the **replyTo** destination specified in the request message. The client deletes the temporary queue after the response has been received. The server is able to handle simultaneous requests from multiple clients since each incoming request message contains the destination to which the reply should be sent.

**When running the WSDLJava command tool on a Web Services Description Language (WSDL) document that contains a JMS-style endpoint URL, for example, jms:/..., the following error displays: java.io.IOException: WSW3099E: Error: Emitter failer. Invalid endpoint address in port <x> in service <y>: <jms-url-string>. How can the problem be resolved?**

This error most likely occurred because the `urlprotocols.jar` file was not in the CLASSPATH when the `WSDL2Java` command ran. The `urlprotocols.jar` file contains the custom protocol handler for the JMS protocol, and it must be in the CLASSPATH when the `WSDL2Java` command runs, if your WSDL document contains a JMS URL. To resolve this problem, edit the `<WAS_HOME>/bin/setupCmdLine.bat` (Windows) or `<WAS_HOME>/bin/setupCmdLine.sh` (UNIX) file, and locate the line which sets the `WAS_CLASSPATH` environment variable. Add `%WAS_HOME%/lib/urlprotocols.jar` (Windows) or `$WAS_HOME/lib/urlprotocols.jar` (UNIX) to the end of that line. Make sure to use the proper delimiter character for your platform, a ";" (semi-colon) for Windows platforms and a ":" (colon) for UNIX platforms.

### **The malformed URL exception displays when running a client that uses a JMS endpoint URL. How can the problem be resolved?**

If you are running a managed client with the `launchClient` command, edit the `install_root\bin\setupCmdLine.bat` file on Windows platforms or `install_root/bin/setupCmdLine.sh` file on UNIX platforms, so that the `WAS_CLASSPATH` environment variable contains the `$WAS_HOME/lib/urlprotocols.jar` file for UNIX platforms or `%WAS_HOME%\lib\urlprotocols.jar` on Windows platforms. See the answer to the question "When running the WSDLJava command tool on a Web Services Description Language (WSDL) document that contains a JMS-style endpoint URL..." for more details. If you are running an unmanaged client, make sure that the `$WAS_HOME/lib/urlprotocols.jar` file on UNIX platforms, or the `%WAS_HOME%\lib\urlprotocols.jar` file on Windows platforms, is in the CLASSPATH when you run your client.

---

## **Web services: Resources for learning**

Use the following links to find relevant supplemental information about getting started with Web services. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas. The following sections are covered in this reference:

View links to additional information about:

- Web services overview: Purpose, planning and designing to use Web services
- Developing Web services Java API for XML-based remote procedure call (JAX-RPC) and the J2EE programming model
- Security
- Administration
- Samples
- Other references

### **Web services overview: Purpose, planning and designing to use Web services**

- ✓ **IBM Web Services architecture debuts**  
(<http://www.ibm.com/developerworks/webservices/library/w-int.html?dwzone=webservices>)

Introducing IBM Web services, a distributed software architecture of service components. This brief overview and in-depth interview on IBM DeveloperWorks cover the fundamental concepts of Web services architecture and what they mean for developers. The interview with IBM professional Rod Smith explores which types of developers Web services targets, how Web services reduces development time, what developers could be doing with Web services now, and takes a glance at the economics of dynamically discoverable services.

- ✓ **Web services (r)evolution, Part 1** (<http://www-106.ibm.com/developerworks/library/ws-peer1.html>)  
This article focuses on the benefits and challenges of building Web services applications. Web services might be an evolutionary step in designing distributed applications, however, they are not without their problems. Outlined are the difficulties developers face in creating a truly workable distributed system of Web services. This article also outlines author Grahm Glass' plan for building peer-to-peer Web applications.










### Developing Web services

- ✓ **JSR 109: Implementing Enterprise Web Services** (<http://jcp.org/en/jsr/detail?id=109>)  
This document describes the J2EE specification model.
- ✓ **Java API for XML-based RPC (JAX-RPC): Core Web Services API in the Java platform** (<http://java.sun.com/xml/jaxrpc/>)  
This document reviews the JAX-RPC which enables Java technology developers to develop SOAP based interoperable and portable Web services.
- ✎ **SOAP** (<http://www.w3.org/TR/SOAP>)  
This article is a detailed overview of SOAP, which includes programming specifications.
- ✎ **Web Services Description Language** (<http://www.w3.org/TR/wsdl>)  
This article is a detailed overview of Web Services Description Language (WSDL), which includes programming specifications.
- ✎ **Universal Description, Discovery and Integration** (<http://www.uddi.org/about.html>)  
This article is a detailed overview of Universal Description, Discovery and Integration (UDDI).
- ✎ **UDDI4J: Matchmaking for Web services** (<http://www-106.ibm.com/developerworks/library/ws-uddi4j>)  
Reviewed in this article are the basics of UDDI, the Java API to UDDI, and how you can use this technology to start building, testing, and deploying your own Web services.

### Security

- ✎ **Security in a Web Services World: A Proposed Architecture and Roadmap** (<http://www-106.ibm.com/developerworks/webservices/library/ws-secmap/>)  
This document describes a proposed model for addressing security within a Web service environment. It defines a comprehensive Web Services Security model that supports, integrates, and unifies several popular security models, mechanisms, and technologies, including both symmetric and public key technologies, in a way that enables a variety of systems to securely interoperate

in a platform and language-neutral manner. It also describes a set of specifications and scenarios that show how these specifications can be used together.

-  **Web Services Security (WS-Security)** (<http://www-106.ibm.com/developerworks/library/ws-secure/>)  
The Web Services Security specifications describe enhancements to SOAP messaging to provide quality of protection through message integrity, message confidentiality, and single message authentication. These mechanisms can be used to accommodate a wide variety of security models and encryption technologies. Web Services Security also provides a general-purpose mechanism for associating security tokens with messages. Additionally, Web Services Security describes how to encode binary security tokens. Specifically, the specification describes how to encode X.509 certificates and Kerberos tickets, as well as how to include opaque encrypted keys. It also includes extensibility mechanisms that can be used to further describe the characteristics of the credentials that are included with a message.
-  **Web Services Security Addendum** (<http://www-106.ibm.com/developerworks/library/ws-secureadd.html>)  
This document describes clarifications, enhancements, best practices, and errata of the Web Services Security specification.
-  **WS-Security Profile of the OASIS Security Assertion Markup Language (SAML) Working Draft 04, 10 September 2002** (<http://www.oasis-open.org/committees/security/docs/draft-sstc-ws-sec-profile-04.pdf>)  
This document proposes a set of standards for SOAP extensions used to increase message confidentiality.
-  **Web Services Security: Soap Message Security Working Draft 12, Monday 21 April 2003** (<http://www.oasis-open.org/committees/download.php/1686/WSS-SOAPMessageSecurity-12-04021.pdf>)  
This document describes the support for multiple token formats, trust domains, signature formats, and encryption technologies.
-  **JSR 55:Certification Path API** (<http://jcp.org/en/jsr/detail?id=55>)  
This document provides a short description of the certification path API.
-  **XML-Signature Syntax and Processing** (<http://www.w3.org/TR/xmlsig-core/>)  
This document specifies XML digital signature processing rules and syntax. XML signatures provide integrity, message authentication, or signer authentication services for data of any type, whether located within the XML that includes the signature or elsewhere.
-  **Canonical XML Version 1.0** (<http://www.w3.org/TR/xml-c14n>)  
This specification describes a method for generating a physical representation, the canonical form, of an XML document that accounts for the permissible changes.
-  **Exclusive XML Canonicalization Version1.0** (<http://www.w3.org/TR/xml-exc-c14n/>)  
Canonical XML [XML-C14N] specifies a standard serialization of XML that, when applied to a subdocument, includes the subdocument's ancestor context including all of the namespace declarations and attributes in the "xml:" namespace.
-  **XML Encryption Syntax and Processing** (<http://www.w3.org/TR/xmlenc-core/>)

This document specifies a process for encrypting data and representing the result in XML.





-  **Decryption Transform for XML Signature**

(<http://www.w3.org/TR/xmlenc-decrypt>)

This document specifies an XML Signature "decryption transform" that enables XML Signature applications to distinguish between those XML Encryption structures that were encrypted before signing, and must not be decrypted, and those that were encrypted after signing, and must be decrypted, for the signature to validate.

-  **WS-Security** (<http://schemas.xmlsoap.org/ws/2002/04/secext/>)

This document specifies resources for the April 2002 Web Services Security Specification. The following addendums and drafts are available:

-  <http://schemas.xmlsoap.org/ws/2002/07/secext/>  
(<http://schemas.xmlsoap.org/ws/2002/07/secext/>)
-  <http://schemas.xmlsoap.org/ws/2002/07/utility/>  
(<http://schemas.xmlsoap.org/ws/2002/07/utility/>)
-  **OASIS draft 12 for secext**  
(<http://schemas.xmlsoap.org/ws/2003/06/secext/>)
-  **OASIS draft 12 for utility**  
(<http://schemas.xmlsoap.org/ws/2003/06/utility/>)

#### Administration


-  **SOAP Security Extensions: Digital Signature**

(<http://www.w3.org/TR/SOAP-dsig>)

This document specifies the syntax and processing rules of a SOAP header entry to carry digital signature information within a SOAP 1.1 Envelope

-  **Apache Software Foundation** (<http://www.apache.org>)


#### Samples

- Two SOAP samples are available. The samples include information about implementing SOAP services and SOAP security. Refer to the  **Samples Gallery** (<http://www7b.software.ibm.com/wsdd/library/samples/AppServer.html>) for these samples.

#### Other references

-  **Web services insider, Part 1: Reflections on SOAP** (<http://www-106.ibm.com/developerworks/webservices/library/ws-ref1>)

What is the current state of the *Web services revolution*? Find out at this Web site that features the column *Web services insider, Part 1*. The author answers this question by reviewing the tools and technologies that have emerged over the past year, highlighting their differences and similarities.

-  **The Web services insider, Part 2: A summary of the W3C Web Services Workshop** (<http://www-106.ibm.com/developerworks/webservices/library/ws-ref2>)

This is a brief summary of a W3C Web services workshop.

---

## Web services implementation scope

Use this page to view and manage the scope of the ports of a Web Service bean.



To view this administrative console page, click **Applications >Enterprise Applications > *application\_instance* > Web Module > *module\_instance* >Web Services Implementation Scope**.

## Port

Specifies a port name for a Web service. A module can contain one or more Web services, each of which can contain one or more ports.

## Service

Specifies the name of the Web service.

## URI

Specifies the Uniform Resource Identifier (URI) of the binding file that defines the scope. The URI is relative to the Web module.

## Scope

Specifies the scope of a port.

The scope determines when a new instance of a service implementation is created for the Web service ports in a module. An application scope causes the same instance of the implementation to be used for all requests on the application. A session scope causes the same instance to be used for all requests on each session. A request scope causes a new instance to be used on every request.

---

## Web services client bindings

Use this page to specify the Web Service Descriptive Language (WSDL) file name and default port type mappings for the Web services in a module.

To view this page, click **Applications >Enterprise Applications > *application\_instance* > Web Module > *module\_instance* >Web Services Client Bindings**.

For EJB modules, click **Applications >Enterprise Applications > *application\_instance* > EJB Module > *module\_instance* >Web Services Client Bindings**

## Web Service

Specifies the name of this Web service.

## URI

Specifies the Uniform Resource Identifier (URI) of the binding file that defines the scope. The URI is relative to the module.

## WSDL Filename

Specifies the the WSDL file name, which is relative to the module.

## Default Port Mappings

Specifies and manages the default port type mapping for a Web service.

---

## Default Port Mapping Definitions collection

Use this page to view and manage a default port type mapping for a Web service.

To view this page of the Administrative Console, click **Applications** > **Enterprise Application** > *application\_instance* > **Web Module** > *module\_instance* > **Web Services Client Bindings** > **Edit** > *default\_port\_instance*.

For EJB modules, click **Applications** > **Enterprise Application** > *application\_instance* > **EJB Module** > *module\_instance* > **Web Services Client Bindings** > **Edit** > *default\_port\_instance*.

Specify the default port of a service when a particular port type is requested. The port type is described by its local name and namespace. A `getPort` method specifying only the port type gets the port named by the default port local name and namespace.

### Port Type Local Name

Specifies the name of this Web service.

### Port Type Namespace

Specifies the local name describing the port type to be mapped.

### Default Port Local Name

Specifies the namespace describing the port type to be mapped.

### Default Port Namespace

Specifies the namespace of the port to map to.

---

## Default Port Type Mapping Properties settings

Use this page to view and manage a default port type mapping for a Web service.

To view this page of the Administrative Console, click **Applications** > **Enterprise Application** > *application\_instance* > **Web Module** > *module\_instance* > **Web Services Client Bindings** > **Edit** > *default\_port\_instance*.

For EJB modules, click **Applications** > **Enterprise Application** > *application\_instance* > **EJB Module** > *module\_instance* > **Web Services Client Bindings** > **Edit** > *default\_port\_instance*.

Specify the default port of a service when a particular port type is requested. The port type is described by its local name and namespace. A `getPort` method specifying only the port type gets the port named by the default port local name and namespace.

### Port Type Local Name

Specifies the name of this Web service.

### Port Type Namespace

Specifies the local name describing the port type to be mapped.

## Default Port Local Name

Specifies the namespace describing the port type to be mapped.

## Default Port Namespace

Specifies the namespace of the port to map to.

---

## Publish WSDL files settings

Use this page to publish Web Services Descriptive Language (WSDL) files.

To view this page, click **Applications** > **Enterprise Applications** > *application\_instance* > **Publish WSDL**.

When you click **OK**, a zip file of the application's Web services-enabled modules is produced. The name of the published zip file is *application\_name\_WSDLFiles.zip*. In the published zip file, the directory structure is *application\_name/module\_name/[META-INF|WEB-INF]/wsdl/wsdl\_file\_name*.

In a published WSDL file, the location attribute of a service soap:address stanza contains the URL through which the Web service is accessed. You can specify the portion of the URL to be used for the Web services in each module. You can access the Web services in a module through a HTTP transport or JMS transport, or both. You can specify URL information for both types of transports.

## HTTP

Specifies the the *protocol* (either http or https), *host\_name*, and *port\_number* to be used in the URL.

The URL prefix format is *protocol://host\_name:port\_number*, for example, *http://myHost:9045*. The actual URL that appears in a published WSDL file consists of the prefix prepended to the module's context-root and the Web service url-pattern, for example, *http://myHost:9045/services/myService* .

## Select HTTP URL prefix

Specifies the drop down list associated with a default list of URL prefixes. This list is the intersection of the set of ports for the module's virtual host and the set of ports for the module's application server. Use items from this list if the Web services application server is accessed directly.

## Custom HTTP URL prefix

Specifies the *protocol*, *host*, and *port\_number* of the intermediate service if the Web services in a module are accessed through an intermediate node, for example the Web Services Gateway or an IHS server.

## JMS

Specifies the JMS URL prefix string used for each module.

The URL prefix specified must contain the destination and connectionFactory properties. It can contain other property-value pairs, but it must not contain the targetService property, which is added by the system when the published WSDL files are created. The format of the JMS URL prefix is *jms:[queue|topic]?destination=target\_queue\_or\_topic|connectionFactory=factory\_name*, for example, *jms:/queue?destination=jms/Q1|connectionFactory=jms/QCF*. The actual

URL that appears in a published WSDL file consists of the prefix prepended to the Web service targetService, for example, `jms:/queue?destination=jms/Q1|connectionFactory=jms/QCF|targetService=StockQuote`.

---

## Using Apache SOAP Web services in Version 5.0 and 5.0.1

### Before you begin

This topic and the topics linking to it describe how to use Apache SOAP 2.3 support for WebSphere Application Server Version 5.0 and 5.01. This support will be deprecated in a future release. If you are planning a new Web services project, it is recommended that you use Web services for J2EE.

Version 5.02 introduces support for Web services for J2EE, which replaces the "IBM WebSphere Web Services for J2EE Technology Preview" available as a separate download from <http://www7b.boulder.ibm.com/wsdd/downloads/techpreviews.html>. This supports emerging Java Web services standards like JAX-RPC and Web services for J2EE. It is recommended that new development efforts use the and follow these standards.

You can find more information about working with Web services by visiting the Internet sites listed in [Web services: Resources for learning](#).

Developing and managing Apache Simple Object Access Protocol (SOAP) 2.3-based Web services explains how to work with Web services that are written directly to Apache SOAP, Version 2.3 API. To develop and manage an Apache SOAP 2.3-based Web service:

### Steps for this task

1. Develop an Apache SOAP client.
2. Assemble Apache SOAP Web services.
3. Deploy Apache SOAP Web services applications.
4. Administer deployed Apache SOAP Web services (SOAP-XML administrative tool).
5. Deploy Web services.
6. Secure Apache SOAP services.

## Developing an Apache SOAP client

### Before you begin

Develop a client that has a Web Services Description Language (WSDL) file or service implementation.

Creating clients to access the Simple Object Access Protocol (SOAP) services published in WebSphere Application Server is a straightforward process. The Apache SOAP implementation, integrated with WebSphere Application Server, contains a client API to assist in SOAP client application development.

The SOAP API documentation is available in [WebSphere Application Server Java documentation](#).

To create a client that interacts with a SOAP Remote Procedure Call (RPC) service:

### Steps for this task

1. Obtain the interface description of the SOAP service.  
This description provides the signatures of the methods that you want to invoke. You can either look at a Web Services Description Language (WSDL) file for the service, or view the service itself to see its implementation.
2. Create the call object.  
The SOAP call object is the main interface to the underlying SOAP RPC code.
3. Set the target Uniform Resource Identifier (URI) in the call object using the `setTargetObjectURI()` method.  
Pass the Uniform Resource Name (URN), a type of URI, that the service uses for its identifier, in the deployment descriptor.
4. Set the method name that you want to invoke in the call object using the `setMethodName()` method.  
This method must be one of the methods exposed by the service located at the URN from the previous step.
5. Create the necessary parameter objects for the RPC call and then set them in the call object using the `setParams()` method.  
Ensure you have the same number and same type of parameters as those required by the service.
6. Execute the call object `invoke()` method and retrieve the response object.  
Remember the RPC call is synchronous, so it can take some time to complete.
7. Check the response for a fault using the `getFault()` method, and then extract any results or returned parameters.  
While most of the providers only return a result, the DB2 stored procedure provider can also return output parameters.

### What to do next

Interacting with a document-oriented SOAP service requires you to use lower-level Apache SOAP API calls. You must first construct an envelope object containing the contents of the message, including the body and any headers, that you wish to send. Then create a message object where you invoke the `send()` method to perform the actual transmission.

To create a secure SOAP service, do the following:

1. Create a simple object.
2. Define an envelope editor.
3. Specify a pluggable envelope editor.
4. Define the transports.

Your code can look like the following example:

```
EnvelopeEditor editor=  
new PluggableEnvelopeEditor(new InputSource(conf), home);  
SOAPTransport transport =  
new FilterTransport(editor, new SOAPHTTPConnection());  
call.setSOAPTransport(transport);
```

The characteristics of the secure session are specified by the `conf` configuration file.

## Accessing enterprise beans with Apache SOAP

Using Apache Simple Object Access Protocol (SOAP) to call enterprise beans is done in the same manner as calling Java bean methods using Apache SOAP.

A Web service can be a simple stateless session bean that performs number processing and returns a data value. When the client code makes a call to the data processing method of this service and an instance of the stateless session is not available, the Apache SOAP run time does the following:

- Calls the enterprise bean create method to obtain a stateless session
- Calls the requested method

## Assembling Apache SOAP Web services

Use the Simple Object Access Protocol (SOAP) EAR Enabler tool to assemble Apache SOAP Web services. The following tasks comprise assembling Web services:

### Steps for this task

1. Create or locate the software resource to be exposed as a service.  
To expose the service, create a programming artifact, one of the supported types, or locate an existing piece of code of the supported type.
2. (Archive (EAR) file).  
Package the code artifact into an EAR file. This step is a deployment packaging requirement of WebSphere Application Server. Use the Application Assembly Tool (AAT) to package the artifact.
3. Create the Apache SOAP deployment descriptor for the desired service.  
In order to deploy an artifact as a SOAP service, create a Apache SOAP deployment descriptor that describes the service you are creating. This step exposes the programming artifact as a service. The descriptor describes and defines the parts of the code that will be invoked with the SOAP calls.  
The information contained in the deployment descriptor varies, depending on the type of artifact you are exposing. For more information about various deployment descriptor types see the article Simple Object Access Protocol deployment descriptor. For example, the following deployment descriptor might be used with the *StockQuoteSample*:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn" [type="message"]>
  <isd:provider type="java"
    scope="Request | Session | Application"
    methods="exposed-methods">
    <isd:java class="implementing-class" [static="true|false"]/>
  </isd:provider>
<isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>
```

4. Execute the SoapEarEnabler tool to enable your Web service.  
Your code artifact must first be packaged into an EAR file. Next, using the deployment descriptor as input, add the necessary pieces to the EAR file to enable the artifact as a Web service. To facilitate this process, use the Java based tool called SoapEarEnabler. Depending on whether you secure the Web service, this tool will add two Web modules: soap.war and soap-sec.war to the EAR file. These Web modules include the SOAP deployment descriptors, plus the necessary parts to deploy the service into the WebSphere Application Server run time.

The service does not become available until the soap-enabled EAR file is installed, and the server is restarted.

## Apache SOAP deployment descriptor

Apache SOAP utilizes XML documents called deployment descriptors to provide the SOAP run-time with information on client services.

Deployment descriptors provide a variety of information including:

- Service Uniform Resource Name (URN) used to route the request when it arrives
- Method and class details, if the service is being provided by a Java class
- User ID and password information, if the service provider must connect to a database

There are four types of deployment descriptors that you can use in WebSphere Application Server.

- Standard Java class
- Enterprise bean
- Bean Scripting Framework (BSF)
- DB2 stored

The contents of the deployment descriptor depend on the type of artifact that is being exposed using SOAP. Samples of these deployment descriptors are available in the SOAP 2.3 sample. You can link to the sample from the article [Web services: Resources for Learning](#).

### Standard Java class deployment descriptor

A deployment descriptor that exposes a service implemented with a standard Java class, including a normal Java bean. An example of the standard Java class deployment descriptor follows:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn" [type="message"]>
  <isd:provider type="java"
    scope="Request | Session | Application"
    methods="exposed-methods">
    <isd:java class="implementing-class" [static="true|false"]/>
  </isd:provider>
<isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>
```

where:

- *service-urn* is the uniform resource name (URN) given to a service. All services deployed within a single EAR file must have URNs that are unique within that EAR file.
- *exposed-methods* is a list of methods, separated by spaces, which are being exposed.
- *implementing-class* is a fully-qualified class name (that is, a *packagename.classname*) that provides the methods that are being exposed.
- *type* is an optional attribute on the service <element>. Type is set to the value "message" if the service is document-oriented instead of RPC-invoked.

- *static* is an optional attribute called on the `java>` element, which may be set to either "true" or "false", depending on whether the methods are exposed or not exposed. If exposed, this attribute indicates whether the method is static or not static.
- *scope* is an attribute on the `<provider>` element, which indicates the lifetime of the instantiation of the implementing class.
- "Request" indicates the object is removed after the request completes.
- "Session" indicates the object lasts for the current lifetime of the HTTP session.
- "Application" indicates the object lasts until the servlet, which is servicing the requests, is terminated.

## EJB deployment descriptor

A deployment descriptor that exposes a service implemented with an Enterprise Java Bean. An example of an EJB deployment descriptor is as follows:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn">
  <isd:provider type="provider-class"
    scope="Application"
    methods="exposed-methods">
    <isd:option key="JNDIName" value="jndi-name"/>
    <isd:option key="FullHomeInterfaceName" value="home-name" />
  </isd:provider>
  <isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>
```

where:

- *service-urn* and *exposed-methods* have the same meaning as in the standard Java class deployment descriptor
- *provider-class* with a stateless session bean follows:  
com.ibm.soap.providers.WASStatelessEJBProvider
- *jndi-name* is the registered JNDI name of the enterprise bean
- *home-name* is the fully-qualified class name of the enterprise bean's home.

The default values for the iiop URL and context provider keys are:

```
<isd:option key="ContextProviderURL" value="iiop://localhost:2809" />
<isd:option key="FullContextFactoryName" value="com.ibm.websphere.naming.WsnInitialContextFactory" />
```

To use your own values, you must specify:

```
<isd:option key="ContextProviderURL" value="<URL to the JNDI provider>" />
<isd:option key="FullContextFactoryName" value="<Context factory full class name>" />
```

## Bean Scripting Framework (BSF) script deployment descriptor

A deployment descriptor that exposes a service implemented with a BSF script. An example follows:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn">
  <isd:provider type="script"
    scope="Request | Session | Application"
    methods="exposed-methods">
    <isd:script language="language-name" [source="source-filename"]>[script-body]
  </isd:script>
  </isd:provider>
  <isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>
```



where:

- *service-urn*, *exposed-methods*, and *scope* have the same meaning as in the standard Java class deployment descriptor.
- *language-name* is the name of the BSF-supported language that is used to write the script.
- *source* or *script-body* is an attribute that is mandatory for the <script> element. The *script-body* attribute contains the actual script that is used to provide the service.
- *source-filename* refers to the file which contains the service implementation if the deployment descriptor has the *source* attribute.

### DB2 stored procedure deployment descriptor

A deployment descriptor that exposes one or more DB2 stored procedures as a service. An example follows:

```
<isd:service xmlns:isd="http://xml.apache.org/xml-soap/deployment"
  id="urn:service-urn">
  <isd:provider type="com.ibm.soap.providers.WASDB2SPPProvider"
    scope="Application"
    methods="* | exposed-methods">
    <isd:option key="userID" value="db-userid"/>
    <isd:option key="password" value="db-password"/>
    [<isd:option key="fullContextFactoryName" value="context-factory"/>]
    <isd:option key="datasourceJNDI" value="jndi-name"/>]
    [<isd:option key="dbDriver" value="db-driver"/>]
    <isd:option key="dbURL" value="db-url"/>]
  </isd:provider>
  <isd:faultListener>org.apache.soap.server.DOMFaultListener</isd:faultListener>
</isd:service>
```

where:

- *service-urn* and *exposed-methods* have the same meaning as in the standard Java class deployment descriptor.
- *db-userid* is a valid user ID used to access the database where the stored procedures reside.
- *db-password* is a valid password for the specified user ID  
The *db-userid* and *db-password* entries are optional. These entries can be set in the datasource. In WebSphere Application Server, the preferred way for administering the *db-userid* and *db-password* entries is with a datasource. Changing the user ID and password is easier when the information is located in a datasource rather than in a separate deployment descriptor file.
- *context-factory* is the name of the context factory used to access the database
- *jndi-name* is the datasource used to access the database
- *db-driver* is the database driver used to access the database.  
A *db-driver* is not required if a data source JNDI name is specified.
- *db-url* is the URL that specifies the database to access.

The *methods* attribute on the <provider> element can contain a list of space-separated procedure names to expose, or an \* (asterisk). An asterisk indicates all available stored procedures should be exposed.

## Enabling Apache SOAP Web services in an enterprise application

The Simple Object Access Protocol (SOAP) EAR enabler tool is a Java application that enables a set of SOAP services within an enterprise archive (EAR) file.

The SoapEAREnabler (SEE) tool guides you through the required steps to enable one or more services within an application. The tool makes a backup copy of your original EAR file in the event you need to remove or add services at a later time.

Before invoking the SEE tool, create an SOAP deployment descriptor for each service to be enabled.

Follow these steps to enable a Web service:

### Steps for this task

1. Invoke the SEE tool from `was_install/bin` directory using one of the following commands:
  - Use the command **SoapEarEnabler** on the Windows platform
  - Use the command **SoapEarEnabler.sh** on UNIX platforms.
2. Start with an existing EAR file, either one created with the Application Assembly Tool (AAT), or a previously-created, valid, J2EE-compliant EAR file. The SEE tool will **not** accept a SOAP-enabled EAR file as input.
3. After you enable Web services, install the EAR file in WebSphere Application Server.
4. You will be prompted as to whether you wish to add the administration client to the EAR. This is a Web-based client that will allow you to list all active services for a specific context from a browser window. With this interface, you can stop and start existing services. You might choose to not add this interface for security reasons, or you might want to secure the interface before making a service available.

### What to do next

You have invoked the SEE from the WebSphere Application Server bin directory and now you should know how the tool operates and its prompts and commands.

The tool operates in two modes:

1. Interactive
2. Silent

Specify all required command line arguments to use the tool in **silent** mode.

In **interactive** mode, the SEE tool prompts all required information. The following dialog is an example of using the tool in interactive mode.

**Note:** In this dialog, user input is in *italics*, and tool output is in **bold**.

```
SoapEarEnabler (On Windows NT)
SoapEarEnabler.sh (On UNIX platforms)
```

```
Please enter the name of your ear file:
  ..\work\stockquote.ear
```

```
How many services would you like your application
to contain (1...n)?
```

1

Now prompting for info for service #1:

Please enter the file name of the SOAP deployment descriptor xml file:

..\work\StockQuoteDD.xml

Is this service an EJB (y=yes / n=no)?

n

If there are additional jar files that need to be added to the classpath, specify the number to add (answer 0 for none) (0...n)?

1

Classpath addition 1: Here are all the archive files found, please choose one to add to the classpath: ([1] samples.jar, [2] stockquote.war):

1

Should this service be secured ("y/n" versus "y= yes, n= no)?

n

Please enter a context root for your non-secured services (e.g. /soap):

/soapsamples

Do you wish to install the administration client?

Warning! You should not install this client in a production ear unless you intend to secure the URI to it.

Install the administration client ("y/n" versus "y= yes, n= no)?

y

In **Silent mode**, supply the arguments in the same order as for the interactive prompts.

**Note:** In silent mode, the SOAP administrative GUI will not install. Also, you will not be prompted for the SOAP administrative GUI.

The following example describes how to use the tool in silent mode:

**soapearenabler [args]**

where the arguments must be specified in the following order:

**<ear-file-name>**

**<number-of-services>**

*The following block is repeated based on the number of services specified.*

**<deployment-descriptor-file-name>**

**<service-is-an-ejb- (y/n)>**

*The following argument should be supplied only if service-is-an-ejb-(y/n) is y.*

**<ejb-jar-file-uri-(already-in-ear)>**

**<number-of-additional-jar-files**

**(0, 1, 2...)>**

*The following argument is repeated number-of-additional-jar-files times.*

**<classpath-entry-uri-(already-in-ear)>**

**<secure-this-service-(y/n)>**

*This following argument is supplied only if secure-this-service-(y/n) is n for any service.*

**<context-root-for-non-secured-services,**

**ex: /soap>**

*This following argument is supplied only*

*if secure-this-service-(y/n) is y for any service.*

**<context-root-for-secured-services, ex:  
/soapsec>**

## Examples

Silent mode example

The following is an example of deploying one enterprise bean as a non-secured service:

```
soapearenabler soap.ear 1
    d:\xml-soap\java\samples\ejbadder\deploymentdescriptor.xml
y adderservice-ejb.jar 1 samples.jar n /soap
```

The following is an example of deploying one enterprise bean as a non-secured service, and one Java class as a secured service:

```
soapearenabler soap.ear 2
    d:\xml-soap\java\samples\ejbadder\deploymentdescriptor.xml
y adderservice-ejb.jar 1 samples.jar n
    d:\xml-soap\java\samples\stockquote\deploymentdescriptor.xml
n 1 samples.jar y /soap /soap-sec
```

Deploying two Java classes as non-secured services example

```
soapearenabler soap.ear 2
    d:\xml-soap\java\samples\stockquote\deploymentdescriptor.xml
n 1 samples.jar n
    d:\xml-soap\java\samples\addressbook\deploymentdescriptor.xml
n 1 samples.jar n /soap
```

The line breaks in the above examples have been modified for this documentation. Typically, commands are issued on a single line.

## Deploying Apache SOAP Web services applications

To deploy a programming artifact as a Simple Object Access Protocol (SOAP)-accessible Web service in WebSphere Application Server:

### Steps for this task

1. Install the service-enabled EAR file.  
Use the SOAP-enabled EAR file that you created and install it in WebSphere Application Server.
2. Update the Web server plug-in configuration.  
Run the GenPluginCfg.bat file on Windows NT, or the GenPluginCfg.sh script on UNIX platforms, to regenerate the plug-in configuration.
3. Restart the application server.

## Administering deployed Apache SOAP Web services (XML-SOAP administrative tool)

Administer the deployed Apache SOAP Web service with the XML-SOAP administrative tool.

Use the SOAPEarEnabler tool to add administrative interfaces to your EAR files. Then, use the XML-SOAP administrative tool with these EAR files to do the following tasks for each context root:

- List configured services, showing active and stopped services

- Stop a service
- Start a service
- View the Apache SOAP deployment descriptor for a service

Access the XML-SOAP administrative tool through a Web browser by specifying:  
`http://localhost/<contextroot>/admin/index.html`

where *context root* is the context specified when installing the SOAP-enabled .ear file.

### What to do next

You cannot use the XML-SOAP administrative tool to add or remove a service. Use the SoapEarEnabler tool to add or remove services. A stopped service is persisted across starts and stops of the application server. Therefore, if you stop a service, it remains stopped until the next time you use the XML-SOAP administrative tool to start it again.

You can add the XML-SOAP administrative tool interface to an enterprise application when you SOAP-enable the EAR file. In interactive mode, you are asked whether you want to add the XML-SOAP administrative tool interface. Replying *yes* will add the necessary Java Server Pages (JSP) files and bindings allowing access to the XML-SOAP administrative tool interface for the application. The interface is an optional addition because you might not want to expose it in a production environment. Optionally, secure the XML-SOAP administrative tool using the Application Assembly Tool (AAT) to assign roles.

## Securing Apache SOAP Web services

The Simple Object Access Protocol (SOAP) security extension, included with WebSphere Application Server, is a security architecture based on the Apache SOAP security specification, and widely-accepted security technologies such as Secure Sockets Layer (SSL).

There are three options for security when using HTTP as the transport protocol:

- HTTP basic authentication
- SSL (HTTPS)
- SOAP signature

You can combine these security options according to your specific security requirements.

Because the SOAP specification does not address all security issues, several proposals evolved to bridge the security gaps. See *Web services: Resources for Learning* for more information about additional documentation.

### Migrating Apache SOAP security

You can use the following steps when migrating a WebSphere Application Server, Version 4.0 application that uses Apache SOAP 2.2 security, to a WebSphere Application Server, Version 5.0 application that uses Apache SOAP 2.3 security.

#### Steps for this task

1. If you have any scripts or .bat files that set up the CLASSPATH before running your application, make sure that the `j2ee.jarfile`, located in the WebSphere `V5<WebSphere/AppServer>/install root` directory is in the CLASSPATH.

2. (Optional) Replace security key files, if needed.

You might need new security key files if you get the following error:

```
"conf/cl-sig-config.xml" javax.security.cert.CertPathBuilderException:  
No end-entity certificate matching the selection criteria could be found
```

The three security key files are named SOAPclient, SOAPserver, and sslserver.p12. To locate the key files:

- a. If the application is already installed these files are found in the directory <WebSphere/AppServer>/installedApps/<yourAppName>.ear/soapsec.war/key/
- b. If the application is not yet installed, these files are found by expanding <yourAppName.ear>. Then, expand the soapsec.war. The files are in the key/directory.

3. Stop and restart the server if you made any updates or replaced the key files.

### Securing Apache SOAP services with HTTP basic authentication

Many applications require users to provide identifying information. You cannot provide access control for individual services. You can only provide access control for the router servlets, for example, the rpcrouter servlet URI. If you can get to a servlet, you can access any of the Web services served through the servlet. Therefore, if you have a set of secure services, you have to partition them differently so that they are accessed through a URI that is secured, for example, /secureRPCRouter. An example of a service that is not secure or accessible to everyone is /unprotectedRCPRouter.

Using the (Application Assembly Tool (AAT)), you can set authorization levels by assigning roles to HTTP methods and by assigning users to roles. You can then *authenticate* users, verifying they are authorized to view specific information. There are many ways to prompt users for authentication data.

### Securing Apache SOAP services on Secured Sockets Layer

To make a request over HTTPS using the Secured Socket Layer support of Apache SOAP, you need a separate Java Secure Socket Extension (JSSE) provider.

WebSphere Application Server includes the `ibmjsse.jar` in the IBM Developer Kit for Windows, Java Technology Edition.

The *SOAP on SSL* scenario is useful for many business-to-business (B2B) applications because:

- The data in transit is protected from eavesdropping or forgery by SSL.
- The client identity is authenticated through user ID and password, which are encrypted by the SSL transport.

For example, if an inventory application is configured as a Web service, the service provider has the following two SOAP service entries:

- `https://foo.com/inventory/inquiry`
- `https://foo.com/inventory/update`

Each SOAP service entry should be deployed as a separate enterprise application (EAR) because each service has a different access control policy, which is: anyone can inquire about the inventory but only the inventory clerks can update the contents.

The SOAP enablement model limits you to one context root for the unsecured services and another for the secured services. In this example, you want to make the inquiry service unsecured and the update service secured. If you want different levels of security for a secured service, then you must deploy the entries in the secured service as separate EAR files.

To enable the *SOAP on SSL* scenario:

#### Steps for this task

1. Configure the Web server (`httpd.conf`) so that it only allows SSL access to these servlets.
2. Configure the security role for the `RPCRouterServlet` in the inquiry services EAR file.

By doing this step, the `RPCRouterServlet` for the *inquiry* service is accessible by everyone, while the `RPCRouterServlet` for the *update* service requires authentication based on the HTTP basic authentication (userID and password).

In this case, the *update* application does not know the identity of the requester; it only knows that access is granted. In other words, the *update* application is not concerned with the identity of the user because it knows WebSphere Application Server is ensuring that only authenticated users have access.

### **Securing Apache SOAP services on Secured Sockets Layer with SOAP Signature**

Applications might need non-repudiable proof of exchanged messages. One example is a Web service that accepts part orders. The business partners establish a form of trust relationship based on public keys. This can be done using the public key infrastructure (PKI) through a third party certificate authority (CA), or by exchanging public keys with a secure channel. The following service is deployed with a signature verification function:

`https://foo.com/partorder`

Configure signature verification with the following information:

- Scope of signature (indicates the portion of the Simple Object Access Protocol (SOAP) envelope that must be authenticated. The default is the content of *SOAP-ENV:Body*).
- Trusted keys or trusted root keys.
- Default key to verify signature if no `KeyInfo` is specified.
- Other policies regarding signature validation.
- Behavior when signature verification fails.
- Additional requirements on signature (as for example, specific requirements on hash/C14N algorithms to be used, timestamp validity, and so forth).

If the signature is missing or if signature verification fails, the signature verification function can be configured so that the servlet returns a SOAP fault.

To send part orders to the `https://foo.com/partorder` service, the service requester should sign his SOAP messages with a signature component. The signature component is initialized using two templates:

1. `<ds:SignedInfo>` template
2. `<ds:KeyInfo>` template

The `<ds:SignedInfo>` template controls the following:

- What parts of the SOAP envelope must be signed

- What algorithms (canonicalization, transformation, digest, sign) should be used

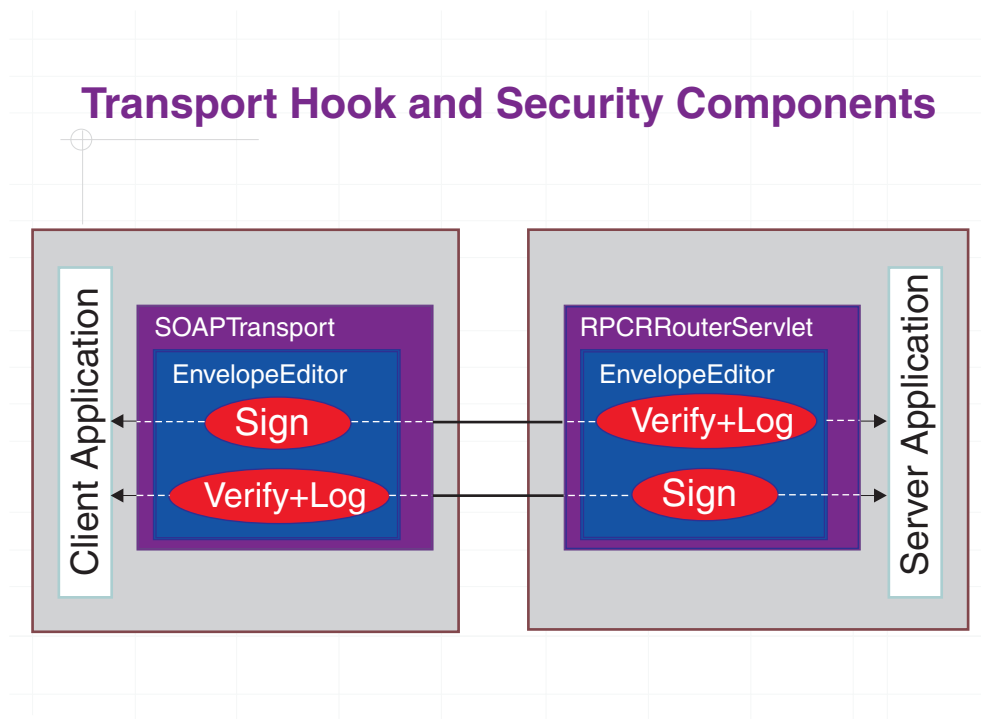
The <ds:KeyInfo> template controls the following:

- Whether or not to include the entire certificate chain in <ds:KeyInfo>
- Decision to include only certificate and serial number
- Public key value
- Decision to provide no key information (so that the default key must be used for verification).

You can combine the service request with HTTP basic authentication, if necessary.

### Apache SOAP signature architecture

An overview of the Apache SOAP signature architecture is illustrated in the figure below.



Using the SOAP transport hook, you can plug-in the following security components:

- Signer
- Verifier, with logging capability

The transport hook is called the *EnvelopeEditor*. A *PluggableEnvelopeEditor* is also provided, which allows you to plug in your security components. As illustrated,



the *EnvelopeEditor* is encapsulated in the *SOAPTransport* on the client side. On the server side, *EnvelopeEditor* is encapsulated in *RPC/MessageRouterServlet*. The same components can be used on either side.

When a client application sends a request, the request is signed and transmitted to the server. At the server side, the request is verified and delivered to a server application or, in the case of a Remote Procedure Call (RPC), to a Java object. The response is processed in the same manner. The verifier component also has a logging function to log the verified messages in a file. Signatures and verifier components are configurable. You can specify encryption, digest message algorithm, certificate path policy, and other security technologies.

For more information about *EnvelopeEditor*, Signature header handler and Verification header handler, see *Web services: Resources for Learning*.

## UDDI4J specifications

The following considerations are specific to the support for UDDI4J provided by WebSphere Application Server, Version 5:

- **UDDI4J class libraries provided.**

WebSphere Application Server provides two UDDI4J class libraries:

- **uddi4jv2.jar**

- This class library contains classes which support Version 2 of the UDDI specification.

- **uddi4j.jar**

- This class library is provided for compatibility with WebSphere Application Server, Version 5, and supports Version 1 of the UDDI specification. The classes in this library are deprecated.

- **UDDI4J error handling**

- When invoking UDDIProxy inquiry methods, UDDIException is thrown if errors are received from the UDDI proxy. UDDIException can contain a DispositionReport with information about the error. APIs that do not return a data object, provide the disposition report. SOAPException is thrown if a communication error occurs or if the resulting data cannot be parsed as a valid SOAP message.

See the article *Web services: Resources for Learning*, to learn more about the use of Simple Object Access Protocol (SOAP), Universal Description, Discovery and Integration (UDDI) Registry, UDDI4J, programming specifications and examples.

## Web services: Resources for learning

Use the following links to find relevant supplemental information about getting started with Web services. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas. The following sections are covered in this reference:

View links to additional information about:

- [Web services overview: Purpose, planning and designing to use Web services](#)

- Developing and using Web services applications and clients
- Programming instructions and examples
- Administration
- Samples
- Other references




### Web services overview: Purpose, planning and designing to use Web services

- ✓ **IBM Web Services architecture debuts**  
(<http://www.ibm.com/developerworks/webservices/library/w-int.html?dwzone=webservices>)  
Introducing IBM Web services, a distributed software architecture of service components. This brief overview and in-depth interview on IBM DeveloperWorks cover the fundamental concepts of Web services architecture and what they mean for developers. The interview with IBM professional Rod Smith explores which types of developers Web services targets, how Web services reduces development time, what developers could be doing with Web services now, and takes a glance at the economics of dynamically discoverable services.
- ✓ **Web services (r)evolution, Part 1** (<http://www-106.ibm.com/developerworks/library/ws-peer1.html>)  
This article focuses on the benefits and challenges of building Web services applications. Web services might be an evolutionary step in designing distributed applications, however, they are not without their problems. Outlined are the difficulties developers face in creating a truly workable distributed system of Web services. This article also outlines author Graham Glass' plan for building peer-to-peer Web applications.




### Developing and using Web services applications and clients

- 📄 **SOAP** (<http://www.w3.org/TR/SOAP>)  
This article is a detailed overview of SOAP, which includes programming specifications.
- ✓ **Building a SOAP client** (<http://www-106.ibm.com/developerworks/library/x-soapcl/index>)  
This article describes a simple, general purpose SOAP client in Java that uses no specialized SOAP libraries. Instead of creating the SOAP request XML document for you under the hood, this client lets you create your own request with any XML or text editor. Instead of merely giving you the remote method return values, the client shows you the actual SOAP response XML document. The short Java program shows exactly what SOAP is all about: opening up an HTTP connection, sending the appropriate XML to invoke a remote method, and then reading the XML response returned by the server.
- 📄 **Web Services Description Language** (<http://www.w3.org/TR/wsdl>) This article is a detailed overview of Web Services Description Language (WSDL), which includes programming specifications.
- 📄 **Universal Description, Discovery and Integration**  
(<http://www.uddi.org/about.html>) This article is a detailed overview of Universal Description, Discovery and Integration (UDDI).
- 📄 **UDDI4J: Matchmaking for Web services** (<http://www-106.ibm.com/developerworks/library/ws-uddi4j>) Reviewed in this article are the basics of UDDI, the Java API to UDDI, and how you can use this technology to start building, testing, and deploying your own Web services.


### Programming instructions and examples

-  **Web Services Description Language** (<http://www.w3.org/TR/wsdl>) This article is a detailed overview of Web Services Description Language (WSDL), which includes programming specifications.
-  **Universal Description, Discovery and Integration** (<http://www.uddi.org/about.html>) This article is a detailed overview of Universal Description, Discovery and Integration (UDDI).
-  **UDDI4J: Matchmaking for Web services** (<http://www-106.ibm.com/developerworks/library/ws-uddi4j>) Reviewed in this article are the basics of UDDI, the Java API to UDDI, and how you can use this technology to start building, testing, and deploying your own Web services.



### Administration

-  **SOAP Security Extensions: Digital Signature** (<http://www.w3.org/TR/SOAP-dsig>) This document specifies the syntax and processing rules of a SOAP header entry to carry digital signature information within a SOAP 1.1 Envelope
-  **Simple Object Access Protocol (SOAP) 1.1** (<http://www.w3.org/TR/SOAP-dsig>) This site offers detailed information about SOAP Signature, including Envelope Editor, Signature Header Handler and Verification Header Handler.
-  **Apache Software Foundation** (<http://www.apache.org>)

### Samples

- Two SOAP samples are available. The samples include information about implementing SOAP services and SOAP security. Refer to the  **Samples Gallery** (<http://www7b.software.ibm.com/wsdd/library/samples/AppServer.html>) for these samples.

### Other references

-  **Web services insider, Part 1: Reflections on SOAP** (<http://www-106.ibm.com/developerworks/webservices/library/ws-ref1>) What is the current state of the *Web services revolution*? Find out at this Web site that features the column *Web services insider, Part 1*. The author answers this question by reviewing the tools and technologies that have emerged over the past year, highlighting their differences and similarities.
-  **The Web services insider, Part 2: A summary of the W3C Web Services Workshop** (<http://www-106.ibm.com/developerworks/webservices/library/ws-ref2>) This is a brief summary of a W3C Web services Workshop.



---

## Chapter 9. Web Services Invocation Framework (WSIF): Enabling Web services

Use this topic to learn how to invoke Web services to use the Web Services Invocation Framework

The Web Services Invocation Framework (WSIF) is a WSDL-oriented Java API that allows you to invoke Web services dynamically, regardless of what format (for example EJB) the service is implemented in, or what mechanism (for example JMS) is used to access it.

WSIF enables you, as a Web services developer, to move away from the usual Web services programming model of working directly with the SOAP APIs, towards a model where you interact with representations of the services. So you can work with the same programming model regardless of how the service is implemented and accessed.

If you want to know more about the issues that WSIF addresses, see [Goals of WSIF](#).

If you want to know how WSIF addresses these issues, see [An overview of WSIF](#).

To use WSIF, see the following topics:

- [Using WSIF to invoke Web services](#).
- [WSIF system management and administration](#).
- [WSIF API](#).

For more information about working with WSIF, visit the Internet sites listed in [WSIF: Resources for Learning](#).

---

### Goals of WSIF

SOAP bindings for Web services are part of the WSDL specification. So when most developers think of using a Web service, they immediately think of assembling a SOAP message and sending it across the network to the service endpoint, using some SOAP client API. For example: with Apache SOAP the client creates and populates a Call object which encapsulates the service endpoint, the identification of the SOAP operation to be invoked, the parameters that have to be sent, and so on.

While this works for SOAP, it is limited in its use as a general model for invoking Web services for the following reasons:

- Web services are not just SOAP services.
- Tying client code to a particular protocol implementation is restricting.
- Incorporating new bindings into client code is hard.
- Multiple bindings can be used in flexible ways.
- A freer Web services environment enables intermediaries.

The goals of WSIF are therefore:

- To give a binding-independent mechanism for Web service invocation.

- To free client code from the complexities of any particular protocol used to access a Web service.
- To enable dynamic selection between multiple bindings to a Web service.
- To help the development of Web service intermediaries.

## **WSIF - Web services are not just SOAP services**

You can deploy as a Web service any program with a WSDL description of its functional aspects and access protocols; and in the J2EE environment, the same component is available over multiple transports and protocols.

For example, you can have a database stored procedure, which is then exposed as a stateless session bean, and then deployed into a SOAP router to become a SOAP service. At each stage, the fundamental service is the same. All that changes is the access mechanism: from JDBC to RMI-IIOP and then to SOAP.

The WSDL specification defines a SOAP binding for Web services, but you can add binding extensions to the WSDL so that, for example, you can offer an enterprise bean as a Web service using RMI/IIOP as the access protocol. You can even treat a single Java class as a Web service, with in-thread Java method invocations as the access protocol. With this broader definition of a Web service, you need a binding-independent mechanism for service invocation.

## **WSIF - tying client code to a particular protocol implementation is restricting**

If your client code is tightly bound to a client library for a particular protocol implementation, it can become hard to maintain. For example if you move from Apache SOAP to JMS or enterprise bean, the process can take a lot of time and effort. To avoid these problems, you need a protocol implementation-independent mechanism for service invocation.

## **WSIF - incorporating new bindings into client code is hard**

As is explained in Web services are not just SOAP services, if you want to make an application that uses a custom protocol work as a Web service, you can add extensibility elements to WSDL to define the new bindings. But in practice, achieving this is hard. For example you have to design the client APIs for using this protocol; and if your application uses just the abstract interface of the Web service, you have to write tools to generate the stubs that enable an abstraction layer. These are tasks that can take a lot of time and effort. What you need is a service invocation mechanism that allows bindings to be updated or new bindings to be plugged in easily.

## **WSIF - multiple bindings can be used in flexible ways**

Imagine that you have successfully deployed an application that uses a Web service offering multiple bindings. For example, imagine that you have a SOAP binding for the service and a local Java binding that lets you to treat the local service implementation (a Java class) as a Web service.

The local Java binding for the service can only be used if the client is deployed in the same environment as the service itself, and if this is the case it is far more efficient to communicate with the service by making direct Java calls than using the SOAP binding.

If your clients could switch the actual binding used based on run-time information, they could choose the most efficient available binding for each situation. In order to take advantage of Web services that offer multiple bindings, you need a service invocation mechanism that allows you to switch between the available service bindings at runtime, without having to generate or recompile a stub.

## **WSIF - a freer Web services environment enables intermediaries**

Web services offer application integrators a loosely-coupled paradigm. In such environments, intermediaries can be very powerful. Intermediaries are applications that intercept the messages that flow between a service requester and a target Web service, and perform some mediating task (for example logging, high-availability or transformation) before passing on the message. They can be as small as a simple Web service, or as large as the Web services gateway. WSIF is designed to make building intermediaries both possible and simple. Using WSIF, intermediaries can add value to the service invocation without needing transport-specific programming.

---

### **An overview of WSIF**

WSIF provides a Java API for invoking Web services, independent of the format of the service or the transport protocol through which it is invoked. It addresses all of the issues identified in the goals of WSIF.

WSIF provides the following features:

- It has an API that provides binding-independent access to any Web service.
- It is closely based on WSDL, so it can invoke any service that can be described in WSDL.
- It allows stubless (completely dynamic) invocation of a Web service.
- You can plug a new or updated implementation of a binding into WSIF at runtime.
- You can defer the choice of a binding until runtime.

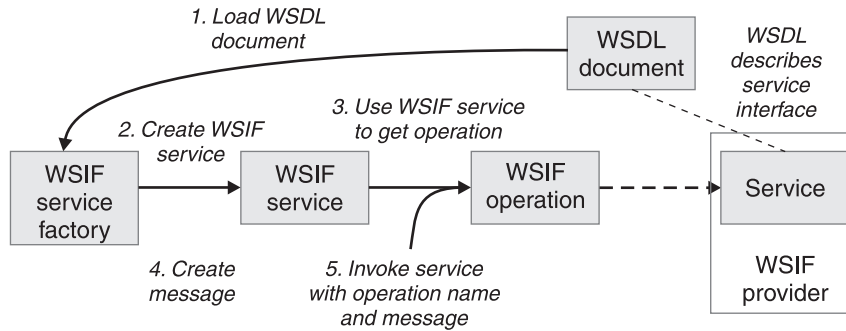
WSIF is designed to work both in an unmanaged environment (stand-alone) and inside a managed container. You can use JNDI to find the WSIF service, or else read in the WSDL definition.

For more conceptual information about WSIF and WSDL, see the following topics:

- WSIF and WSDL
- WSIF architecture
- Using WSIF with Web services that offer multiple bindings
- WSIF usage scenarios
- Dynamic invocation

### **WSIF architecture**

The WSIF architecture is shown in the following figure. The components of this architecture are described after the figure.



### WSIF provider

A WSIF provider is an implementation of a WSDL binding that can run a WSDL operation through a binding-specific protocol. WebSphere Application Server includes WSIF providers for SOAP over HTTP, SOAP over JMS, Java, enterprise beans, and Native JMS. For more information, see *Using the WSIF providers*.

### WSIFOperation

The runtime representation of an operation, called *WSIFOperation* is responsible for invoking a service based on a particular binding. For more information, see *WSIF API reference: Using ports*.

### WSIFService

The *WSIFService* is responsible for generating an instance of *WSIFOperation* to be used for a particular invocation of a service operation. For more information, see *Finding a port factory or service*.

### WSDL documents

The Web service WSDL document contains the location of the Web service. The binding document defines the protocol and format for operations and messages defined by a particular portType.

## Using WSIF with Web services that offer multiple bindings

Using WSIF, a client application can dynamically choose the optimal binding to use to invoke Web service operations.

For example, a Web service provider could offer a SOAP binding for the service and a local Java binding that allows you to treat the local service implementation (a Java class) as a Web service. If the client is deployed in the same environment as the service, the local Java binding for the service can be used and provides more efficient communication with the service by making direct Java calls rather than using the SOAP binding.

For more information on how to configure a client to dynamically select between multiple bindings, see *Developing a WSIF service*.

## WSIF and WSDL

WSDL is the acronym for Web Services Description Language (<http://www.w3.org/TR/wsdl>).

In WSDL a service is defined in three distinct parts:

- The **PortType**. The PortType defines the abstract interface offered by the service. A PortType defines a set of *Operations*. Each operation can be In-Out (request-response), In-Only, Out-Only and Out-In (Solicit-Response). Each



operation defines the input and/or output *Messages*. A message is defined as a set of *Parts* and each part has a schema-defined type.

- The **Binding**. A binding defines how to map between the abstract PortType and a real service format and protocol. For example the SOAP binding defines the encoding style, the SOAPAction header, the namespace of the body (the targetURI), and so on.
- The **Port**. This defines the actual location (endpoint) of the available service - for example, the HTTP URL on which a SOAP service is available.

Currently in WSDL, each Port has one and only one binding, and each binding has a single PortType. But (more importantly) each Service (PortType) can have multiple Ports, each of which represents an alternative location and binding for accessing that service.

WSIF follows the semantics of WSDL as much as possible:

- The WSIF dynamic invocation API directly exposes runtime equivalents of the model from WSDL. For example, invocation of an operation involves executing an "Operation" with an Input Message.
- WSDL has extension points that allow new ports and bindings to be added so that WSDL can describe new systems. The equivalent concept in WSIF is a provider, that allows WSIF to understand a class of extensions and therefore support new service implementation types.

As a metadata-based invocation framework, WSIF follows the design of the metadata. As WSDL is extended, WSIF is updated to follow.

**Note:** The **implicit and primary** type system of WSIF is XML Schema, not Java. WSIF supports invocation using dynamic proxies, which support Java type systems, but when you use the WSIFMessage interface it is your responsibility to populate WSIFMessage objects with data based on the XML Schema types as defined in the WSDL document. So you should define types of objects by a canonical and fixed mapping from Schema types into the runtime.

For more information on WSDL, see WSIF: Resources for learning.

## WSIF usage scenarios

This topic describes two brief scenarios that illustrate the role WSIF plays in the emerging Web services environment.

### Scenario: Redevelopment and redeployment

If you are implementing Web services today you are probably working with simple prototypes. As your Web services move into production, you will need to reimplement and redeploy them. WSIF uses the same API calls with different underlying technologies. So if you use WSIF you can reimplement and redeploy your services without changing the client code, and you can use existing highly reliable and high-performance infrastructures like RMI-IIOP and JMS without sacrificing the location-independence that the Web service model offers.

### Scenario: Service Flow composition

A service flow typically invokes a Web service, then passes the response from one Web service into the next Web service, perhaps performing some transformation in the middle.

There are two key aspects to this that WSIF provides:

- A representation of the service invocation based on the metadata in WSDL.
- The ability to build invocations based on the portType only, which can be used on any implementation.

For example, imagine that you build a "meta-service" that uses a number of services to build a process. Initially several of those services are simple JavaBean prototypes that are written and exposed through SOAP, but you plan to reimplement some of them as EJB components, and to out-source others.

If you use SOAP, it ties up multiple threads for every onward invocation, as they pass through the webserver and servlet engine into the SOAP router. If you use WSIF to call the beans directly, you get much better performance compared to SOAP and you don't lose access or location transparency. Using WSIF, you can move the JavaBean implementations to EJB implementations without changing the client code, and to move some of the Web services from local implementations to external SOAP services you just update the WSDL.

## Dynamic invocation

In WSIF, dynamic invocation means providing the following levels of support when invoking Web services:

1. Support for WSDL extensions and bindings that were not known at build time.
2. Support for Web services that were not known at build time.

WSIF supports (1) above through the use of providers.

The providers support (2) above by using the WSDL description to access the target service.

---

## Using WSIF to invoke Web services

You invoke a Web service dynamically by using the WSIF API directly. You only specify the location of the WSDL file for the service, the name of the operation to be invoked, and any operation arguments needed. All the information needed to access the Web service is available through WSDL; the abstract interface, the binding, and the service endpoint.

This kind of invocation does not require stub classes and does not need a separate compilation cycle.

More information on using WSIF to invoke Web services is given in the following topics:

- Using the WSIF providers.
- Developing a WSIF service.
- Using complex types.
- Using JNDI.
- **5.0.2** Passing SOAP messages with attachments using WSIF.
- Interacting with the WebSphere J2EE container.
- Running WSIF as a client.

## Using the WSIF providers

A WSIF provider is an implementation of a WSDL binding that can run a WSDL operation through a binding-specific protocol.

Providers implement the interface between the WSIF API and the actual implementation of a service. Providers are pluggable within the WSIF framework, and are registered based upon the namespace of the WSDL extension that they implement.

WebSphere Application Server includes the following WSIF providers:

- SOAP (over HTTP) provider.
- JMS providers (SOAP over JMS, and Native JMS).
- Java provider.
- EJB provider.

### Note:

- Some providers use the J2EE programming model to utilize J2EE services.
- If a provider is available, but its required class libraries are not, the provider is disabled.

## Using the SOAP provider

The SOAP provider allows WSIF stubs and dynamic clients to invoke SOAP services. The provider supports SOAP 1.1 (<http://www.w3.org/TR/SOAP/>) over HTTP. The WSIF SOAP Provider utilizes Apache SOAP 2.3 for parsing and creating SOAP messages, but is not limited to invoking services served by Apache SOAP.

The WSIF SOAP provider supports:

- SOAP-ENC encoding.
- **5.0+** **5.0.1+** RPC style SOAP messages.
- **5.0.2** RPC style and Document style SOAP messages.
- **5.0.2** SOAP messages with attachments.

The SOAP provider is not transactional.

**Note:** Before you deploy to WebSphere Application Server a Web service that you expect to be used by multiple clients connecting over SOAP, you must set up your application's deployment descriptor file (dds.xml) to handle multiple connections correctly. For more information see WSIF troubleshooting tips.

### What to do next

For an example of the sort of code changes that need to be made in the WSDL file for a SOAP provider, see the following topics:

- The SOAP over JMS provider - writing the WSDL extension.
- **5.0.2** SOAP messages with attachments - writing the WSDL extensions.

## Using the JMS providers

JMS is an API for transport technology. The mapping to a JMS destination is defined during deployment and maintained by the container.

The JMS Destination Web service endpoint may be realized in any of the following ways:

- The JMS destination for the queue can be the Web service implementation.
- The JMS destination can be (but is not required to be) associated with a Message Driven Bean by the EJB container, thereby allowing the Message Driven Bean to be the Web service implementation.
- (For SOAP over JMS) The JMS Destination can unwrap the JMS message and route the SOAP message to a stateless session bean Web service implementation.

The JMS Destination must respect the interaction model expected by the client and defined by the WSDL. It must return a response if one is required.

When the JMS Destination creates the JMS response message the following rules must be adhered to:

- The response message must be sent to the **JMSReplyTo** from the incoming request.
- The **JMSCorrelationID** of the response message must be set to the value of the **JMSMessageID** from the request message.
- The response must be sent with a **deliveryMode** equal to the **JMSDeliveryMode** of the request message.
- The response must be sent with a **priority** equal to the **JMSPriority** of the request message.
- The **timetolive/JMSExpiration** must be set a value equal to the **JMSExpiration** of the request message.

The client does not see any of these headers. The container receives the JMS message and (for SOAP over JMS) removes the SOAP message to send to the client.

#### What to do next

See also the following topics:

- Using the SOAP over JMS provider
- Using the native JMS provider
- (The JMS providers - configuring the client and server)

### **Using the SOAP over JMS provider**

#### Before you begin

For information on working with the JMS API, see Using the JMS providers.

The SOAP message, including the SOAP envelope, is wrapped with a JMS message and put on the appropriate Queue. The container receives the JMS message and removes the SOAP message to send to the client.

#### What to do next

For detailed implementation information, see the following topics:

- The SOAP over JMS provider - writing the WSDL extension
- (The JMS providers - configuring the client and server)

## The SOAP over JMS provider - writing the WSDL extension

SOAP protocol messages are carried on the JMS transport with the JMS body type `TextMessage` if the message is strictly XML.

### Usage scenario

#### WSDL Extensions

#### WSDL Binding

The WSDL Binding for SOAP over JMS varies only slightly from the SOAP over HTTP binding. The transport element under the `soap:binding` indicates that JMS is being used.

**RPC style:** If the style is set to "rpc", then it is assumed that an operation is being invoked on the Web service endpoint. The Java parameters and response holders are encoded in the same way as for the standard WSIF SOAP binding:

```
<soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/jms"/>
```

When a SOAP/JMS binding is being used, the `<wsdl:port>` must contain a `<jms:address>` element to identify the JMS Queue to be used.

#### Address:

The address within the port within the service within the WSDL provides the information required for a client to correctly connect to the Web service using the JMS programming model. Typically, it will be the stubs generated to support the SOAP/JMS binding that will act as the JMS client. This does not preclude the Web service client from using the JMS programming model directly. The address element under the service and port must take this form:

```
<jms:address
  destinationStyle="queue"
  jmsVendorURI="http://ibm.com/ns/mqseries"?
  initialContextFactory="com.ibm.NamingFactory"?
  jndiProviderURL="iiop://something:900/wherever"?
  jndiConnectionFactoryName="orange"
  jndiDestinationName="fred"
/>
```

where "?" means "optional".

The vendor URI is a string which uniquely identifies the JMS implementation. WSIF ignores this URI, which is used by the client developer and perhaps the client implementation to determine if it has access to the correct JMS provider in the client runtime.

The `connectionFactory` attribute gives the name of a JMS `ConnectionFactory` object, which can be looked up within the JNDI context given by the `jndiContext` attribute. This `ConnectionFactory` can be used to create a JMS connection to the JMS provider instance that owns the Queue. In a simple configuration this `ConnectionFactory` is the actual `ConnectionFactory` used by the server message listener, and by the clients. However both server and clients can use different `ConnectionFactories`, provided that they all create Connections to the same JMS provider instance.

The JNDI usage pattern is this:

```
<jms:address destinationStyle="queue"
  jndiConnectionFactoryName="orange"
  jndiDestinationName="fred"
>
```

This usage pattern requires the runtime to have a default JNDI provider configured. The provider URL and context factory can also be added to this:

```
<jms:address destinationStyle="queue"
  initialContextFactory="com.ibm.Naming"
  jndiProviderURL="iiop://server:900/"
  jndiConnectionFactoryName="orange"
  jndiDestinationName="fred"
>
```

### Headers and Properties

JMS headers and properties can be set using the `<jms:property>` extension within the binding. This maps a part of a message into a JMS property. For example:

```
<jms:property name="Priority" {part="requestPriority" | value="fixedValue"}/>
```

The header can also be set in the `jms:address` component with a literal value:

```
<jms:property name="Priority" value="fixedValue"/>
```

This binding extension is shared with the Native JMS binding.

### WSDL Binding example

Here is an example of a WSDL that defines a SOAP over JMS binding:

```
<!-- Example: SOAP over JMS Text Message -->

<?xml version="1.0" encoding="UTF-8"?>
<wsdl:definitions
  name="StockQuoteInterfaceDefinitions"
  targetNamespace="urn:StockQuoteInterface"
  xmlns:tns="urn:StockQuoteInterface"
  xmlns:xsd="http://www.w3.org/2000/10/XMLSchema"
  xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
  xmlns:wsdl="http://schemas.xmlsoap.org/wsdl/">
  <wsdl:message name="GetQuoteInput">
    <part name="symbol" type="xsd:string"/>
  </wsdl:message>
  <wsdl:message name="GetQuoteOutput">
    <part name="value" type="xsd:float"/>
  </wsdl:message>

  <wsdl:portType name="StockQuoteInterface">
    <wsdl:operation name="GetQuote">
      <wsdl:input message="tns:GetQuoteInput"/>
      <wsdl:output message="tns:GetQuoteOutput"/>
    </wsdl:operation>
  </wsdl:portType>

  <wsdl:binding name="StockQuoteSoapJMSBinding" type="tns:StockQuoteInterface">
    <soap:binding style="rpc"
      transport="http://schemas.xmlsoap.org/soap/jms"/>
    <wsdl:operation name="GetQuote">
      <soap:operation soapAction="urn:StockQuoteInterface#GetQuote"/>
      <wsdl:input>
```

```

        <soap:body use="encoded" namespace="urn:StockQuoteService"
            encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
    </wsdl:input>
    <wsdl:output>
        <soap:body use="encoded" namespace="urn:StockQuoteService"
            encodingStyle="http://schemas.xmlsoap.org/soap/encoding/" />
    </wsdl:output>
</wsdl:operation>
</wsdl:binding>
<wsdl:service name="StockQuoteService">
    <wsdl:port name="StockQuoteServicePort"
        binding="sqi:StockQuoteSoapJMSBinding">
        <jms:address destinationStyle="queue"
            jndiConnectionFactoryName="myQCF"
            jndiDestinationName="myQ"
            initialContextFactory="com.ibm.NamingFactory"
            jndiProviderURL="iiop://something:900/" />
    </wsdl:port>
</wsdl:service>
</wsdl:definitions>

```

## Using the native JMS provider

### Before you begin

For information on working with the JMS API, see [Using the JMS providers](#).

The native JMS provider allows WSIF clients to treat a JMS destination as a Web service.

### What to do next

For detailed implementation information, see the following topics:

- [The native JMS provider - writing the WSDL extension](#)
- [\(The JMS providers - configuring the client and server\)](#)

## The native JMS provider - writing the WSDL extension

The WSDL extensions for JMS are identified with the namespace prefix "jms". For example <jms:binding>.

### Operations

The supported operations are either One-way operations (send for JMS point-to-point messaging or publish for JMS publish/subscribe messaging) or Request-Response operations (send/receive for JMS point-to-point messaging).

The WSDL operations therefore specify an input message only, or a WSDL input and output message.

### Message Header Data

JMS does not make assumptions about message headers. If the JMS provider is MQSeries, then each JMS message carries an RFH2 header, however, data in this message header is only accessed indirectly by getting and setting JMS message properties.

### Quality of Service Attributes

The Native JMS binding of WSDL operations provides the capability to set message properties. See "JMS Message Header Fields and Properties" below.

## JMS Message Types

JMS message body types are specified in the WSDL binding using the `jms:binding` WSDL extension:

```
<wsdl:binding ... >
  <jms:binding type="messageBodyType"/>
  ...
</wsdl:binding>
```

where `messageBodyType` is either `ObjectMessage` or `TextMessage`.

## Input and Output Messages

Input and output messages are used in the WSDL input and output sections of the WSDL binding, respectively:

```
<wsdl:input ... >
  <jms:input parts="part1 part2 ..."/>
</wsdl:input>

<wsdl:output ... >
  <jms:output parts="part1 part2 ..."/>
</wsdl:output>
```

For JMS Text messages and JMS Object messages created from one or more WSDL message parts, the `jms:input` and `jms:output` WSDL extension specifies the message parts to be used for the JMS message. If no parts are defined, then all message parts will be used. This is used in the WSDL input and output sections of the WSDL binding.

The WSDL message has one part that contains the complete message body. This message body might be the result of a mapping from some other representation (see "Data Mapping" below).

```
<wsdl:input ... >
  <jms:input parts="part1"/>
</wsdl:input>
```

## Fault Messages

Operations that describe message interfaces with a native JMS binding do not have fault messages. No assumptions are made about the message schema or the semantics of message properties, so no distinction can be made between output and fault messages.

## Data Mapping

Mapping of data types is specified using the WSDL format binding extension:

```
<wsdl:binding ... >
  <jms:binding type="..."/>

  <format:typeMapping encoding="Java" style="Java">
    <format:typeMap typeName="..." formatType="targetType"/>
  </format:typemapping>
  ...
</wsdl:binding>
```

The value of `targetType` is dependent on the JMS message type (discussed above) as follows:

- For JMS Object messages, the target data type implements `java.io.Serializable`.



- For JMS Text messages, the target data type is always `java.lang.String`.

Note that the format binding is also used in other bindings that deal with Java interfaces.

### JMS Message Header Fields and Properties

You can set JMS Properties in the WSDL. When you want your application to pass a property into WSIF as a part on the WSIF message, use a `<jms:property>`. When you want to hard code an actual property value into the WSDL, use a `<jms:propertyValue>`. You can put `<jms:property>` and `<jms:propertyValue>` in the `<input>` and `<output>` in the binding operation, and in the `<jms:address>`. `<jms:property>` takes precedence over `<jms:propertyValue>` and properties in the binding operation take precedence over properties in the `<jms:address>`.

For the association with WSDL parts, the `<jms:property>` and `<jms:propertyValue>` WSDL extensions are used as shown below:

```
<wsdl:input ... >
    <jms:property name="propertyName" part="partName"/>
    <jms:propertyValue name="propertyName"
        type="xsdType" value="actualValue"/>
</wsdl:input>
<wsdl:output ... >
    <jms:property name="propertyName" part="partName"/>
</wsdl:output>
```

where *propertyName* identifies the JMS property that is associated with the header field, and *partName* identifies the message part that is associated with the property.

The *propertyName* is the name of one of the predefined JMS message header fields or a user-defined property. Possible values and Java types for the predefined message header fields are shown in the following table:

| Value            | Java type                    |
|------------------|------------------------------|
| JMSMessageId     | java.lang.String             |
| JMSTimeStamp     | long                         |
| JMSCorrelationId | byte [ ] or java.lang.String |
| JMSReplyTo       | javax.jms.Destination        |
| JMSDestination   | javax.jms.Destination        |
| JMSDeliveryMode  | int                          |
| JMSRedelivered   | boolean                      |
| JMSType          | java.lang.String             |
| JMSExpiration    | long                         |

See the JMS specification for restrictions that apply for setting JMS header field values. Attempts to set such restricted values are ignored.

For application-defined JMS message properties, the Java types used in the Native JMS binding implementation (used for calls to the corresponding JMS methods) is derived from the XML schema type in the abstract interface (`wsdl:part`) and the type mapping information in the format binding (`format:typemap`).

The **jms:propertyValue** contains a specification of a literal value and its associated XML schema type.

For **wsdl:output**, only **jms:property** can be specified.

## Transactions

Independent of other quality of service attributes, the asynchronous processing of request-response operations has implications for callers running in a transaction scope. The send request part and the receive response part must be separated into **two transactions** as the send needs to be committed in order for the request message to become visible.

So implementations that process WSDL for asynchronous request-response operations (such as WSIF) must take the following additional actions:

- They must return a correlation id to the user, and provide a **callback** that allows users to pass in the response message in order to process the "second half" of the operation.
- (Optionally) They might implement their own response message "listener" in order to recognize the arrival of response messages, and to manage the correlation to the request message.

## Usage scenario

## WSDL Examples

### Example 1: JMS Text Message (Request-Response)

The JMS text message contains a **java.lang.String**. In this example, the WSDL message contains only one part that represents the whole message body.

```
<!-- Example 1: JMS Text Message -->

<wsdl:definitions ... >

  <!-- simple or complex types for input and output message -->
  <wsdl:types> ... </wsdl:types>

  <wsdl:message name="JmsOperationRequest"> ... </wsdl:message>
  <wsdl:message name="JmsOperationResponse"> ... </wsdl:message>

  <wsdl:portType name="JmsPortType">
    <wsdl:operation name="JmsOperation">
      <wsdl:input name="Request"
        message="tns:JmsOperationRequest"/>
      <wsdl:output name="Response"
        message="tns:JmsOperationResponse"/>
    </wsdl:operation>
  </wsdl:portType>

  <wsdl:binding name="JmsBinding" type="JmsPortType">
    <jms:binding type="TextMessage"/>

    <format:typemapping style="Java" encoding="Java">
      <format:typemap name="xsd:String" formatType="String"/>
    </format:typemapping>

    <wsdl:operation name="JmsOperation">
      <wsdl:input message="JmsOperationRequest">
        <jms:input parts="requestMessageBody"/>
      </wsdl:input>
```

```

        <wsdl:output message="JmsOperationResponse">
            <jms:output parts="responseMessageBody"/>
        </wsdl:output>
    </wsdl:operation>
</wsdl:binding>

<wsdl:service name="JmsService">
    <wsdl:port name="JmsPort" binding="JmsBinding">
        <jms:address destinationStyle="queue"
            jndiConnectionFactoryName="myQCF"
            jndiDestinationName="myDestination"/>
    </wsdl:port>
</wsdl:service>
</wsdl:definitions>

```

## Example 2: JMS Message with Accessing Application-Defined JMS Properties (Request-Response)

As an extension of the previous JMS message example, the following WSDL describes a request-response operation where specific JMS property values of the request and response message are set for the request message and retrieved from the response message.

The JMS properties in the request message are set according to the values in the input message. Likewise, selected JMS properties of the response message are copied to the corresponding values of the output message. The direction of the mapping is determined by the appearance of the **jms:property** tag in the input or output section, respectively.

```

<!-- Example 2: JMS Message with JMS Properties -->

<wsdl:definitions ... >

    <!-- simple or complex types for input and output message -->
    <wsdl:types> ... </wsdl:types>

    <wsdl:message name="JmsOperationRequest">
        <wsdl:part name="myInt" type="xsd:int"/>
        ...
    </wsdl:message>

    <wsdl:message name="JmsOperationResponse">
        <wsdl:part name="myString" type="xsd:String"/>
        ...
    </wsdl:message>

    <wsdl:portType name="JmsPortType">
        <wsdl:operation name="JmsOperation">
            <wsdl:input name="Request"
                message="tns:JmsOperationRequest"/>
            <wsdl:output name="Response"
                message="tns:JmsOperationResponse"/>
        </wsdl:operation>
    </wsdl:portType>

    <wsdl:binding name="JmsBinding" type="JmsPortType">
        <!-- the JMS message type may be any of the above -->
        <jms:binding type="..."/>

        <format:typemapping style="Java" encoding="Java">
            <format:typemap name="xsd:int" formatType="int"/>
            ...
        </format:typemapping>

```

```

<wsdl:operation name="JmsOperation">
  <wsdl:input message="JmsOperationRequest">
    <jms:property message="tns:JmsOperationRequest" parts="myInt"/>
    <jms:propertyValue name="myLiteralString"
      type="xsd:string" value="Hello World"/>
    ...
  </wsdl:input>
  <wsdl:output message="JmsOperationResponse">
    <jms:property message="tns:JmsOperationResponse" parts="myString"/>
    ...
  </wsdl:output>
</wsdl:operation>
</wsdl:binding>

<wsdl:service name="JmsService">
  <wsdl:port name="JmsPort" binding="JmsBinding">
    <jms:address destinationStyle="queue"
      jndiConnectionFactoryName="myQCF"
      jndiDestinationName="myDestination"/>
  </wsdl:port>
</wsdl:service>

</wsdl:definitions>

```

## The JMS providers - configuring the client and server

To enable a service to be invoked through JMS by a WSIF client application, complete the following steps:

1. Create a JMS queue for WSIF to send to the request message to.
 

**Note:** This can be done using MQSeries Explorer, or an equivalent tool from your JMS implementation package.
2. Create a JNDI definition for that JMS queue, and a JNDI queue connection factory.
 

**Note:** If you are using MQSeries, this can be done using the jmsadmin tool.
3. Put the JNDI names of the queue and queue connection factory, as well as your JNDI configuration, in the WSDL file.

You should also be aware of the following specific ways in which WSIF interacts with JMS:

- Only input JMS properties are supported.
- WSIF needs two queues when invoking an operation: one for the request message and one for the reply. The replyTo queue is by default a temporary queue which WSIF creates on behalf of the application. You can specify a permanent queue by setting the JMSReplyTo property to the JNDI name of a queue.
- WSIF uses the default values for properties set by the JMS implementation. However in MQSeries and some other JMS implementations, messages are by default persistent, and the default temporary queue is temporary dynamic and so cannot have persistent messages written to it. So your JMS listener may fail to write a persistent response message to the temporary replyTo queue.
 

**Note:** If you are using MQSeries, you need to create a temporary model queue which is permanent dynamic, then pass this model as the tempmodel of your queue connection factory. This will ensure that persistent messages can be written to the permanent dynamic temporary replyTo queue.

## Using the Java provider

The WSIF Java Provider allows WSIF to invoke Java classes and JavaBeans. This means that, in a "thin-client" environment such as a JVM or Tomcat test runtime, you can define "shortcuts" to local Java code.

The WSIF Java Provider is not intended to be used in a J2EE environment. There is a difference between a client using the WSIF Java Provider to invoke a Java component, and implementing a Web service as a Java component on the server side.

The Java binding exploits the format binding for type mapping. The format binding allows WSDL to define the mapping between XML Schema types and Java types.

The Java provider requires the targeted Java classes to be in the class path of the client. The Java method is invoked synchronously, in-process, in-thread, with the current thread and ORB contexts.

The Java provider is not transactional.

#### What to do next

For examples of the code changes that need to be made in the WSDL file, see The Java provider - writing the WSDL extension.

### **The Java provider - writing the WSDL extension**

The Java provider allows the invocation of a method on a local Java object. In order to use the Java provider you require the following binding specified in the WSDL:

#### Usage scenario

```
<!-- Java binding -->
<binding .... >
  <java:binding />
  <format:typeMapping style="Java" encoding="Java"/>?
  <format:typeMap name="qname" formatType="nmtoken"/>*
</format:typeMapping>
<operation>*
  <java:operation
    methodName="nmtoken"
    parameterOrder="nmtoken"
    returnPart="nmtoken"?
    methodType="instance|constructor"/>
    <input name="nmtoken"? />?
    <output name="nmtoken"? />?
    <fault name="nmtoken"? />?
  </operation>
</binding>
```

where "?" means "optional" and "\*" means "0 or more".

#### **Note:**

- The `format:typeMap` **name** attribute is a qualified name of a simple or complex type used by one of the Java operations.
- The `format:typeMap` **formatType** attribute is the fully qualified Class name for the Java Class that the element specified by **name** maps to.
- The `java:operation` **methodName** attribute is the name of the method on the Java object that is called by the operation.
- The `java:operation` **parameterOrder** attribute contains a whitespace-separated list of part names that define the order in which they are passed to the Java Object's method.

- The java:operation **methodType** attribute must be set to either "instance" or "constructor". The value specifies whether the method being invoked on the object is an instance method or a constructor for the object.

```

<service ... >
  <port>*
    <java:address
      className="nmtoken"/>
  </port>
</service>

```

**Note:** The java:address **className** attribute specifies the fully qualified class name of the object containing the method to invoke.

## Using the EJB provider

The EJB Provider allows WSIF clients to invoke enterprise beans. The EJB client JAR must be available in the client runtime with the current provider. The enterprise bean is invoked using normal EJB invocation methods, using RMI-IIOP, with the current security and transaction contexts. If the EJB provider is invoked within a transaction then the transaction is passed to the onward service and the standard EJB transaction attribute applies.

If there are multiple implementations of the service, it is up to the provider of the service to make sure that they offer the same semantics. For example, in the case of transactionality the bean deployer should specify TX\_REQUIRES\_NEW to force a new transaction.

### What to do next

For examples of the sort of code changes that need to be made in the WSDL file, see The EJB provider - writing the WSDL.

## The EJB provider - writing the WSDL extension

The EJB provider allows the invocation of an enterprise bean through RMI/IIOP. In order to use the EJB provider you require the following binding specified in the WSDL:

### Usage scenario

```

<!-- EJB binding -->
<binding .... >
  <ejb:binding />
  <format:typeMapping style="Java" encoding="Java"/>?
  <format:typeMap name="qname" formatType="nmtoken"/>*
</format:typeMapping>
<operation>*
  <ejb:operation
    methodName="nmtoken"
    parameterOrder="nmtoken"
    returnPart="nmtoken"?
    interface="remote|home"/>
  <input name="nmtoken"? />?
  <output name="nmtoken"? />?
  <fault name="nmtoken"? />?
</operation>
</binding>

```

where "?" means "optional" and "\*" means "0 or more".

**Note:**

- The format:typeMap **name** attribute is a qualified name of a simpleType or complexType used by one of the EJB operations.
- The format:typeMap **formatType** attribute is the fully qualified Class name for the Java Class that the element specified by **name** maps to.
- The ejb:operation **methodName** attribute is the name of the method on the enterprise bean that will be called by the operation.
- The ejb:operation **parameterOrder** attribute contains a whitespace-separated list of part names which define the order in which they are passed to the EJB method.
- The ejb:operation **interface** attribute must be set to either "remote" or "home". The value specifies the interface of the enterprise bean on which the method named by the **method** attribute is accessible.

```

<service ... >
  <port>*
    <ejb:address
      className="nmtoken"
      jndiName="nmtoken"
      initialContextFactory="nmtoken" ?
      jndiProviderURL="nmtoken" ? />
    </port>
  </service>

```

**Note:**

- The ejb:address **className** attribute specifies the fully qualified class name home interface class of the enterprise bean.
- The ejb:address **jndiName** attribute specifies the full JNDI name which is used to look up the enterprise bean.
- The ejb:address **initialContextFactory** attribute is optional and specifies the initial context factory class.
- The ejb:address **jndiProviderURL** attribute is optional and specifies the jndi provider URL

## Developing a WSIF service

A WSIF service is a Web service that uses WSIF. To develop a WSIF service, you first develop the Web service (or use an existing Web service), then develop the WSIF client based on the WSDL document for that Web service.

There are also two pre-built WSIF samples available for download from the Samples Central (<http://www.ibm.com/websphere/developer/library/samples/AppServer.html>) page of the IBM WebSphere Developer Domain Web site. These are as follows:

- The Address Book sample.
- The Stock Quote sample.

For more information on using the pre-built samples, see the documentation that is included in the download package.

To develop a WSIF service, complete the following steps:

### Steps for this task

1. Develop the Web service.

Use Web services tools to discover, create, and publish the Web service. You can develop Java bean, enterprise bean, and URL Web services. You can use Web service tools to create a skeleton Java bean and a sample application from a

WSDL document. For example, an enterprise bean can be offered as a Web service, using RMI/IIOP as the access protocol. Or you can use a Java class as a Web service, with native Java invocations as the access protocol.

You can use the WebSphere Studio Application Developer to create a Web service from a Java bean, as described in its StockQuote service tutorial. The bean that you use in this scenario returns the last trading price from the Internet Web site [www.xmltoday.com](http://www.xmltoday.com) given a stock symbol. Using the Web Service wizard, you generate a binding WSDL document named `StockQuoteService-binding.wsdl` and a service WSDL document named `StockQuoteService-service.wsdl` from the bean `StockQuoteService.java`. You then deploy the Web service to a Web server, generate a client proxy to the Web service, and generate a sample application that accesses the `StockQuoteService` through the client proxy. You test the StockQuote Web service, publish it using the IBM UDDI Explorer, and then discover the StockQuote Web service in the IBM UDDI Test Registry.

## 2. Develop the WSIF client.

The information you need to develop a WSIF client is given in the following topics:

- [Developing the WSIF client - the Address Book sample](#) gives example code to show how you define a Web service in WSDL.
- [Using the WSIF providers](#) describes the available providers, and gives example code of how their WSDL extensions are coded.
- [WSIF API](#) defines the main interfaces that your client uses to support the invocation of Web services defined in WSDL.

**Note:** The Address Book sample is written for synchronous interaction. If you are using a JMS provider, your WSIF client might need to act asynchronously. WSIF provides two main features that meet this requirement:

- A **correlation service** that assigns identifiers to messages so that the request can be matched up with the (eventual) response.
- A **response handler** that picks up the response from the Web service at a later time.

For more information see the WSIF API topic [WSIFOperation - Asynchronous interactions](#) reference.

## Developing the WSIF client - the Address Book sample

The following code fragments show you how to use the WSIF API to invoke the AddressBook sample Web service dynamically.

### Usage scenario

**This is example code for dynamic invocation of the AddressBook sample Web service using WSIF:**

```
try {
    String wsdlLocation="clients/addressbook/AddressBookSample.wsdl";

    // The starting point for any dynamic invocation using wsif is a
    // WSIFServiceFactory. We create ourselves one via the newInstance
    // method.
    WSIFServiceFactory factory = WSIFServiceFactory.newInstance();

    // Once we have a factory, we can use it to create a WSIFService object
    // corresponding to the AddressBookService service in the wsdl file.
    // Note: since we only have one service defined in the wsdl file, we
    // do not need to use the namespace and name of the service and can pass
    // null instead. This also applies to the port type, although values have
```



```

// been used below for illustrative purposes.
WSIFService service = factory.getService(
    wsdlLocation,    // location of the wsdl file
    null,            // service namespace
    null,            // service name
    "http://www.ibm.com/namespace/wsif/samples/ab", // port type namespace
    "AddressBookPT" // port type name
);

// The AddressBook.wsdl file contains the definitions for two complexType
// elements within the schema element. We will now map these complexTypes
// to Java classes. These mappings are used by the Apache SOAP provider
service.mapType(
    new javax.xml.namespace.QName(
        "http://www.ibm.com/namespace/wsif/samples/ab/types",
        "address"),
    Class.forName("com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress"));
service.mapType(
    new javax.xml.namespace.QName(
        "http://www.ibm.com/namespace/wsif/samples/ab/types",
        "phone"),
    Class.forName("com.ibm.www.namespace.wsif.samples.ab.types.WSIFPhone"));
// We now have a WSIFService object. The next step is to create a WSIFPort
// object for the port we wish to use. The getPort(String portName) method
// allows us to generate a WSIFPort from the port name.
WSIFPort port = null;

if (portName != null) {
    port = service.getPort(portName);
}
if (port == null) {
    // If no port name was specified, attempt to create a WSIFPort from
    // the available ports for the port type specified on the service
    port = getPortFromAvailablePortNames(service);
}

// Once we have a WSIFPort, we can create an operation. We are going to execute
// the addEntry operation and therefore we attempt to create a WSIFOperation
// corresponding to it. The addEntry operation is overloaded in the wsdl ie.
// there are two versions of it, each taking different parameters (parts).
// This overloading requires that we specify the input and output message
// names for the operation in the createOperation method so that the correct
// operation can be resolved.
// Since the addEntry operation has no output message, we use null for its name.
WSIFOperation operation =
    port.createOperation("addEntry", "AddEntryWholeNameRequest", null);

// Create messages to use in the execution of the operation. This should
// be done by invoking the createXXXXXMessage methods on the WSIFOperation.
WSIFMessage inputMessage = operation.createInputMessage();
WSIFMessage outputMessage = operation.createOutputMessage();
WSIFMessage faultMessage = operation.createFaultMessage();

// Create a name and address to add to the addressbook
String nameToAdd="Chris P. Bacon";
WSIFAddress addressToAdd =
    new WSIFAddress (1,
        "The Waterfront",
        "Some City",
        "NY",
        47907,
        new WSIFPhone (765, "494", "4900"));

// Add the name and address to the input message
inputMessage.setObjectPart("name", nameToAdd);
inputMessage.setObjectPart("address", addressToAdd);

```

```

// Execute the operation, obtaining a flag to indicate its success
boolean operationSucceeded =
    operation.executeRequestResponseOperation(
        inputMessage,
        outputMessage,
        faultMessage);

if (operationSucceeded) {
    System.out.println("Successfully added name and address to addressbook\n");
} else {
    System.out.println("Failed to add name and address to addressbook");
}

// Start from fresh
operation = null;
inputMessage = null;
outputMessage = null;
faultMessage = null;

// This time we will lookup an address from the addressbook.
// The getAddressFromName operation is not overloaded in the
// wsdl and therefore we can simply specify the operation name
// without any input or output message names.
operation = port.createOperation("getAddressFromName");

// Create the messages
inputMessage = operation.createInputMessage();
outputMessage = operation.createOutputMessage();
faultMessage = operation.createFaultMessage();

// Set the name to find in the addressbook
String nameToLookup="Chris P. Bacon";
inputMessage.setObjectPart("name", nameToLookup);

// Execute the operation
operationSucceeded =
    operation.executeRequestResponseOperation(
        inputMessage,
        outputMessage,
        faultMessage);

if (operationSucceeded) {
    System.out.println("Successful lookup of name '"+nameToLookup+"' in addressbook");

    // We can obtain the address that was found by querying the output message
    WSIFAddress addressFound = (WSIFAddress) outputMessage.getObjectPart("address");
    System.out.println("The address found was:");
    System.out.println(addressFound);
} else {
    System.out.println("Failed to lookup name in addressbook");
}

} catch (Exception e) {
    System.out.println("An exception occurred when running the sample:");
    e.printStackTrace();
}
}

```

**The code above refers to the following sample method:**

```

WSIFPort getPortFromAvailablePortNames(WSIFService service)
    throws WSIFException {
    String portChosen = null;

    // Obtain a list of the available port names for the service
    Iterator it = service.getAvailablePortNames();
    {

```

```

        System.out.println("Available ports for the service are: ");
        while (it.hasNext()) {
            String nextPort = (String) it.next();
            if (portChosen == null)
                portChosen = nextPort;
            System.out.println(" - " + nextPort);
        }
    }
    if (portChosen == null) {
        throw new WSIFException("No ports found for the service!");
    }
    System.out.println("Using port " + portChosen + "\n");

    // An alternative way of specifying the port to use on the service
    // is to use the setPreferredPort method. Once a preferred port has
    // been set on the service, a WSIFPort can be obtained via getPort
    // (no arguments). If a preferred port has not been set and more than
    // one port is available for the port type specified in the WSIFService,
    // an exception is thrown.
    service.setPreferredPort(portChosen);
    WSIFPort port = service.getPort();
    return port;
}

```

### The web service itself uses the following classes:

#### WSIFAddress:

```

public class WSIFAddress implements Serializable {

    //instance variables
    private int streetNum;
    private java.lang.String streetName;
    private java.lang.String city;
    private java.lang.String state;
    private int zip;
    private WSIFPhone phoneNumber;

    //constructors
    public WSIFAddress () { }

    public WSIFAddress (int streetNum,
                        java.lang.String streetName,
                        java.lang.String city,
                        java.lang.String state,
                        int zip,
                        WSIFPhone phoneNumber) {
        this.streetNum = streetNum;
        this.streetName = streetName;
        this.city = city;
        this.state = state;
        this.zip = zip;
        this.phoneNumber = phoneNumber;
    }

    public int getStreetNum() {
        return streetNum;
    }

    public void setStreetNum(int streetNum) {
        this.streetNum = streetNum;
    }

    public java.lang.String getStreetName() {
        return streetName;
    }
}

```

```

public void setStreetName(java.lang.String streetName) {
    this.streetName = streetName;
}

public java.lang.String getCity() {
    return city;
}

public void setCity(java.lang.String city) {
    this.city = city;
}

public java.lang.String getState() {
    return state;
}

public void setState(java.lang.String state) {
    this.state = state;
}

public int getZip() {
    return zip;
}

public void setZip(int zip) {
    this.zip = zip;
}

public WSIFPhone getPhoneNumber() {
    return phoneNumber;
}

public void setPhoneNumber(WSIFPhone phoneNumber) {
    this.phoneNumber = phoneNumber;
}
}

```

#### **WSIFPhone:**

```

public class WSIFPhone implements Serializable {

    //instance variables
    private int areaCode;
    private java.lang.String exchange;
    private java.lang.String number;

    //constructors
    public WSIFPhone () { }

    public WSIFPhone (int areaCode,
                      java.lang.String exchange,
                      java.lang.String number) {
        this.areaCode = areaCode;
        this.exchange = exchange;
        this.number = number;
    }

    public int getAreaCode() {
        return areaCode;
    }

    public void setAreaCode(int areaCode) {
        this.areaCode = areaCode;
    }

    public java.lang.String getExchange() {
        return exchange;
    }
}

```

```

    }

    public void setExchange(java.lang.String exchange) {
        this.exchange = exchange;
    }

    public java.lang.String getNumber() {
        return number;
    }

    public void setNumber(java.lang.String number) {
        this.number = number;
    }
}

```

#### **WSIFAddressBook:**

```

public class WSIFAddressBook {
    private Hashtable name2AddressTable = new Hashtable();

    public WSIFAddressBook() {
    }

    public void addEntry(String name, WSIFAddress address)
    {
        name2AddressTable.put(name, address);
    }

    public void addEntry(String firstName, String lastName, WSIFAddress address)
    {
        name2AddressTable.put(firstName+" "+lastName, address);
    }

    public WSIFAddress getAddressFromName(String name)
    throws IllegalArgumentException
    {
        if (name == null)
        {
            throw new IllegalArgumentException("The name argument must not be " +
                "null.");
        }
        return (WSIFAddress)name2AddressTable.get(name);
    }
}

```

#### **And here's the corresponding WSDL file for the Web service:**

```

<?xml version="1.0" ?>

<definitions targetNamespace="http://www.ibm.com/namespace/wsif/samples/ab"
    xmlns:tns="http://www.ibm.com/namespace/wsif/samples/ab"
    xmlns:typens="http://www.ibm.com/namespace/wsif/samples/ab/types"
    xmlns:xsd="http://www.w3.org/2001/XMLSchema"
    xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
    xmlns:format="http://schemas.xmlsoap.org/wsdl/formatbinding/"
    xmlns:java="http://schemas.xmlsoap.org/wsdl/java/"
    xmlns:ejb="http://schemas.xmlsoap.org/wsdl/ejb/"
    xmlns="http://schemas.xmlsoap.org/wsdl/">

    <types>
        <xsd:schema
            targetNamespace="http://www.ibm.com/namespace/wsif/samples/ab/types"
            xmlns:xsd="http://www.w3.org/2001/XMLSchema">

            <xsd:complexType name="phone">
                <xsd:element name="areaCode" type="xsd:int"/>
            </xsd:complexType>
        </xsd:schema>
    </types>

```

```

        <xsd:element name="exchange" type="xsd:string"/>
        <xsd:element name="number" type="xsd:string"/>
    </xsd:complexType>

    <xsd:complexType name="address">
        <xsd:element name="streetNum" type="xsd:int"/>
        <xsd:element name="streetName" type="xsd:string"/>
        <xsd:element name="city" type="xsd:string"/>
        <xsd:element name="state" type="xsd:string"/>
        <xsd:element name="zip" type="xsd:int"/>
        <xsd:element name="phoneNumber" type="typens:phone"/>
    </xsd:complexType>

</xsd:schema>
</types>

<message name="AddEntryWholeNameRequestMessage">
    <part name="name" type="xsd:string"/>
    <part name="address" type="typens:address"/>
</message>

<message name="AddEntryFirstAndLastNamesRequestMessage">
    <part name="firstName" type="xsd:string"/>
    <part name="lastName" type="xsd:string"/>
    <part name="address" type="typens:address"/>
</message>

<message name="GetAddressFromNameRequestMessage">
    <part name="name" type="xsd:string"/>
</message>

<message name="GetAddressFromNameResponseMessage">
    <part name="address" type="typens:address"/>
</message>

<portType name="AddressBookPT">
    <operation name="addEntry">
        <input name="AddEntryWholeNameRequest"
            message="tns:AddEntryWholeNameRequestMessage"/>
    </operation>
    <operation name="addEntry">
        <input name="AddEntryFirstAndLastNamesRequest"
            message="tns:AddEntryFirstAndLastNamesRequestMessage"/>
    </operation>
    <operation name="getAddressFromName">
        <input name="GetAddressFromNameRequest"
            message="tns:GetAddressFromNameRequestMessage"/>
        <output name="GetAddressFromNameResponse"
            message="tns:GetAddressFromNameResponseMessage"/>
    </operation>
</portType>

<binding name="SOAPHttpBinding" type="tns:AddressBookPT">
    <soap:binding style="rpc"
        transport="http://schemas.xmlsoap.org/soap/http"/>
    <operation name="addEntry">
        <soap:operation soapAction=""/>
        <input name="AddEntryWholeNameRequest">
            <soap:body use="encoded"
                namespace="http://www.ibm.com/namespace/wsif/samples/ab"
                encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
        </input>
    </operation>
    <operation name="addEntry">
        <soap:operation soapAction=""/>
        <input name="AddEntryFirstAndLastNamesRequest">
            <soap:body use="encoded"

```

```

        namespace="http://www.ibm.com/namespace/wsif/samples/ab"
        encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
    </input>
</operation>
<operation name="getAddressFromName">
    <soap:operation soapAction="">
    <input name="GetAddressFromNameRequest">
        <soap:body use="encoded"
            namespace="http://www.ibm.com/namespace/wsif/samples/ab"
            encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
    </input>
    <output name="GetAddressFromNameResponse">
        <soap:body use="encoded"
            namespace="http://www.ibm.com/namespace/wsif/samples/ab"
            encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
    </output>
</operation>
</binding>

<binding name="JavaBinding" type="tns:AddressBookPT">
    <java:binding/>
    <format:typeMapping encoding="Java" style="Java">
        <format:typeMap typeName="typens:address"
            formatType="com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress"/>
        <format:typeMap typeName="xsd:string" formatType="java.lang.String"/>
    </format:typeMapping>
    <operation name="addEntry">
        <java:operation
            methodName="addEntry"
            parameterOrder="name address"
            methodType="instance"/>
        <input name="AddEntryWholeNameRequest"/>
    </operation>
    <operation name="addEntry">
        <java:operation
            methodName="addEntry"
            parameterOrder="firstName lastName address"
            methodType="instance"/>
        <input name="AddEntryFirstAndLastNamesRequest"/>
    </operation>
    <operation name="getAddressFromName">
        <java:operation
            methodName="getAddressFromName"
            parameterOrder="name"
            methodType="instance"
            returnPart="address"/>
        <input name="GetAddressFromNameRequest"/>
        <output name="GetAddressFromNameResponse"/>
    </operation>
</binding>

<binding name="EJBBinding" type="tns:AddressBookPT">
    <ejb:binding/>
    <format:typeMapping encoding="Java" style="Java">
        <format:typeMap typeName="typens:address"
            formatType="com.ibm.www.namespace.wsif.samples.ab.types.WSIFAddress"/>
        <format:typeMap typeName="xsd:string" formatType="java.lang.String"/>
    </format:typeMapping>
    <operation name="addEntry">
        <ejb:operation
            methodName="addEntry"
            parameterOrder="name address"
            interface="remote"/>
        <input name="AddEntryWholeNameRequest"/>
    </operation>
    <operation name="addEntry">
        <ejb:operation

```

```

        methodName="addEntry"
        parameterOrder="firstName lastName address"
        interface="remote"/>
    <input name="AddEntryFirstAndLastNamesRequest"/>
</operation>
<operation name="getAddressFromName">
    <ejb:operation
        methodName="getAddressFromName"
        parameterOrder="name"
        interface="remote"
        returnPart="address"/>
    <input name="GetAddressFromNameRequest"/>
    <output name="GetAddressFromNameResponse"/>
</operation>
</binding>
<service name="AddressBookService">
    <port name="SOAPPort" binding="tns:SOAPHttpBinding">
        <soap:address
            location="http://localhost/wsif/samples/addressbook/soap/servlet/rpcrouter"/>
    </port>
    <port name="JavaPort" binding="tns:JavaBinding">
        <java:address className="services.addressbook.WSIFAddressBook"/>
    </port>
    <port name="EJBPort" binding="tns:EJBBinding">
        <ejb:address className="services.addressbook.ejb.AddressBookHome"
            jndiName="ejb/samples/wsif/AddressBook"
            classLoader="services.addressbook.ejb.AddressBook.ClassLoader"/>
    </port>
</service>
</definitions>

```

## Using complex types

WSIF supports the use of user-defined complex types through the mapping of complex types to Java classes. This mapping can be specified manually or automatically as described in the sections below:

- “Manual mapping of complex types”
- **5.0.2** “Automatic mapping of complex types” on page 531

**Note:** Any calls to the WSIFService `mapType` and `mapPackage` methods used for manual mapping will override any equivalent mapping information that is produced automatically. This helps to maintain backwards compatibility, and also allows for less standard mappings.

### Manual mapping of complex types

The method to use to create these mappings by hand depends on the provider being used. For the Java and EJB providers, the mappings are specified in the wsdl file in the binding element. The syntax for specifying the mapping is as follows:

```

<binding .... >
    <ejb:binding|java:binding/>
        <format:typeMapping style="Java" encoding="Java"/>?
        <format:typeMap name="qname" formatType="nmtoken"/>?*
    </format:typeMapping>
    ...
</binding>

```

where "?" means "optional" and "\*" means "0 or more".

The `format:typeMap name` attribute is a qualified name of a complex type or simple type used by one of the operations.



The `format:typeMap` **formatType** attribute is the fully qualified Class name for the Java Class that the element specified by **name** maps to.

If using the Apache SOAP provider then the mapping of a complex type to a Java Class is specified in the client code through two methods on the `org.apache.wsif.WSIFService` interface:

```
public void mapType(QName elementType, Class javaType)
```

and

```
public void mapPackage(String namespaceURI, String packageName)
```

The **mapType** allows the user to specify a mapping between a WSDL element and method takes a `QName` representing the complex type or simple type and the corresponding Java Class it maps to.

The **mapPackage** method allows the user to specify a more general mapping between a namespace and a Java package. Any custom complex or simple types whose namespace matches that of the mapping will be mapped to a Java Class in the corresponding package. The name of the actual class is derived from the name of the complex type using standard xml to Java naming conventions.

## Automatic mapping of complex types

### 5.0.2

For complex types defined in the WSDL, where a generated bean is used to represent this in Java, the WSIF programming model requires that a call is made to `WSIFService.mapType()`. This call tells WSIF the package and class name of the bean representing the XML schema type that is identified with a `QName`. To make things easier, `WSIFService.mapPackage()` provides a mechanism to specify a wildcard version of this, where any class within a specified package is mapped to the namespace of a `QName`. This is a mechanism for manually mapping an XML schema type to a Java class and back again (one mapping entry provides a bi-directional mapping).

There are many ways of converting a `QName` representing an XML schema type name to a Java package name and class. By default, the `WSIFServiceFactory` does not automatically map types. To enable this feature, set the `WSIF_FEATURE_AUTO_MAP_TYPES` feature on the `WSIFServiceFactory` instance:

```
WSIFServiceFactory factory = WSIFServiceFactory.newInstance();  
factory.setFeature(WSIFConstants.WSIF_FEATURE_AUTO_MAP_TYPES, new Boolean(true));
```

WSIF maps types by converting the XML Schema type `QName` URI part to a package name, and localpart to a class name. It does this using the `WSIFUtils` methods `getPackageNameFromNamespaceURI` and `getJavaClassNameFromXMLName`.

## Using JNDI

This example task shows you how to use WSIF to bind a reference to a Web service, then look up the reference using JNDI.

You access a Web service through information given in the WSDL document for the service. If you don't know where to find the WSDL document for the service, but you know that it has been registered in a UDDI registry, you look it up in the registry. Java programs access java objects and resources in a similar manner, but using a JNDI interface.

The following example shows how, using WSIF, you can bind a reference to a Web service then look up the reference using JNDI.

### Usage scenario

#### **Specifying the argument values for the Web service**

The Web service is represented in WSIF by an instance of the `org.apache.wsif.naming.WSIFServiceRef` class. This simple Referenceable object has the following constructor:

```
public WSIFServiceRef(
    String WSDL,
    String sNS,
    String sName,
    String ptNS,
    String ptName)
{
    [...]
}
```

where

- *WSDL* is the location of the WSDL file containing the definition of the service.
- *sNS* is the full namespace for the service definition (null can be specified if only one service is defined in the WSDL file).
- *sName* is the local name for the service definition (null can be specified if only one service is defined in the WSDL file).
- *ptNS* is the full namespace for the port type within the service that you want to use (null can be specified if only one port type is available for the service).
- *ptName* is the local name for the port type (null can be specified if only one port type is available for the service).

For example, if the WSDL file for the Web service is available from the URL `http://localhost/WSDL/Example.WSDL` and contains the following service and port type definitions -

```
<definitions targetNamespace="http://hostname/namespace/example"
    xmlns:abc="http://hostname/namespace/abc"
[...]
    <portType name="ExamplePT">
        <operation name="exampleOp">
            <input name="exampleInput" message="tns:ExampleInputMsg"/>
        </operation>
    </portType>
[...]
    <service name="abc:ExampleService">
[...]
    </service>
[...]
</definitions>
```

- then you specify the following argument values for `WSIFServiceRef`:

- *WSDL* is `http://localhost/WSDL/Example.WSDL`
- *sNS* is `http://hostname/namespace/abc`
- *sName* is `ExampleService`
- *ptNS* is `http://hostname/namespace/example`
- *ptName* is `ExamplePT`

#### **Binding the service using JNDI**

To bind the service reference in the naming directory using JNDI, you can use the WebSphere Application Server JndiHelper `com.ibm.websphere.naming.JndiHelper` class as follows:

```
[...]
import com.ibm.websphere.naming.JndiHelper;
import org.apache.wsif.naming.*;
[...]
try {
    Context startingContext = new InitialContext();
    WSIFServiceRef ref = new WSIFServiceRef("http://localhost/WSDL/Example.WSDL",
                                           "http://localhost/WSDL/Example.WSDL",
                                           "http://hostname/namespace/abc",
                                           "ExampleService",
                                           "http://hostname/namespace/example",
                                           "ExamplePT");

    JndiHelper.recursiveRebind(startingContext,
                               "myContext/mySubContext/myServiceRef", ref);

}
catch (NamingException e) {
    // Handle error.
}
[...]
```

### Looking up the service using JNDI

The following code fragment shows the lookup of a service using JNDI:

```
[...]
try {
[...]
```

```
    InitialContext ic = new InitialContext();
    WSIFService myService =
        (WSIFService) ic.lookup("myContext/mySubContext/myServiceRef");
[...]
```

```
}
catch (NamingException e) {
    // Handle error.
}
[...]
```

## Passing SOAP messages with attachments using WSIF

The W3C SOAP Messages with Attachments(<http://www.w3.org/TR/SOAP-attachments>) document describes a standard way to associate a SOAP message with one or more attachments in their native format (for example GIF or JPEG) by using a multipart MIME structure for transport. It defines specific usage of the Multipart/Related MIME media type, and rules for the usage of URI references to refer to entities bundled within the MIME package. It thereby outlines a technique for a SOAP 1.1 (<http://www.w3.org/TR/SOAP/>) message to be carried within a MIME multipart/related message in such a way that the SOAP processing rules for a standard SOAP message are not changed.

WSIF supports passing attachments in a MIME message using the SOAP provider. The attachment is a `javax.activation.DataHandler`. The `mime:multipartRelated`, `mime:part` and `mime:content` tags are used to describe the attachment in the WSDL.

For more information, see the following topics:

- SOAP messages with attachments - writing the WSDL extensions.
- SOAP messages with attachments - passing attachments to WSI.
- SOAP messages with attachments - types and type mappings.

**Note:** The following scenarios are not supported:

- Using DIME.
- Passing in `javax.xml.transform.Source` and `javax.mail.internet.MimeMultipart`.
- Using the `mime:mimeXml` WSDL tag.
- Nesting a `mime:multipartRelated` inside a `mime:part`.
- Using types that extend `DataHandler`, `Image`, and so forth.
- Using types that contain `DataHandler`, `Image`, and so forth.
- Using Arrays or Vectors of `DataHandlers`, `Images`, and so forth.
- Using multiple in/out or output attachments.

The MIME headers from the incoming message are not preserved for referenced attachments. The outgoing message contains new MIME headers for `Content-Type`, `Content-Id` and `Content-Transfer-Encoding` that are created by WSIF.

## **SOAP messages with attachments - writing the WSDL extensions**

### Usage scenario

The following WSDL illustrates a simple operation that has one attachment called `attch`:

```
<binding name="MyBinding" type="tns:abc" >
  <soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/http"/>
  <operation name="MyOperation">
    <soap:operation soapAction=""/>
    <input>
      <mime:multipartRelated>
        <mime:part>
          <soap:body use="encoded" namespace="http://mynamespace"
            encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
        </mime:part>
        <mime:part>
          <mime:content part="attch" type="text/html"/>
        </mime:part>
      </mime:multipartRelated>
    </input>
  </operation>
</binding>
```

### **Note:**

- There must be a part (in this example `attch`) on the input message for the operation (in this example `MyOperation`). There can be other input parts to `MyOperation` that are not attachments.
- In the binding input there must either be a `<soap:body` tag or a `<mime:multipartRelated` tag, but not both.
- For MIME messages, the `soap:body` is inside a `mime:part`. There must only be one `mime:part` that contains a `soap:body` in the binding input and that must not contain a `mime:content` as well, because a content type of `text/xml` is assumed for the `soap:body`.
- There can be multiple attachments in a MIME message, each described by a `mime:part`.
- Each `mime:part` (that is not a `soap:body`) contains a `mime:content` that describes the attachment itself. The `type` attribute inside the `mime:content` is not checked or used by WSIF. It is there to suggest to the application using WSIF what the attachment contains. Multiple `mime:contents` inside a single `mime:part` means

that the backend service will expect a single attachment with a type specified by one of the `mime:contents` inside that `mime:part`.

- The `parts="..."` attribute (optional) inside the `soap:body` is assumed to contain the names of all the MIME parts as well as the names of all the SOAP parts in the message.

## SOAP messages with attachments - passing attachments to WSIF

### Usage scenario

The following code fragment could invoke the service described by the example WSDL given in writing the WSDL extensions:

```
import javax.activation.DataHandler;
. . .
DataHandler dh = new DataHandler(new FileDataSource("myimage.jpg"));
WSIFServiceFactory factory = WSIFServiceFactory.newInstance();
WSIFService service =
    factory.getService("my.wsdl",null,null,"http://mynamespace","abc");
WSIFOperation op = service.getPort().createOperation("MyOperation");
WSIFMessage in = op.createInputMessage();
in.setObjectPart("attch",dh);
op.executeInputOnlyOperation(in);
```

The associated type mapping in the `DeploymentDescriptor.xml` file depends upon your SOAP server. For example if you use Tomcat with SOAP 2.3, then `DeploymentDescriptor.xml` contains the following type mapping:

```
<isd:mappings>
<isd:map encodingStyle="http://schemas.xmlsoap.org/soap/encoding/"
  xmlns:x="http://mynamespace"
  qname="x:datahandler"
  javaType="javax.activation.DataHandler"
  java2XMLClassName="org.apache.soap.encoding.soapenc.MimePartSerializer"
  xml2JavaClassName="org.apache.soap.encoding.soapenc.MimePartSerializer"/>
</isd:mappings>
```

In this case, the backend service is invoked with the following signature:

```
public void MyOperation(DataHandler dh);
```

Attachments can also be passed in to WSIF using stubs:

```
DataHandler dh = new DataHandler(new FileDataSource("myimage.jpg"));
WSIFServiceFactory factory = WSIFServiceFactory.newInstance();
WSIFService service =
    factory.getService("my.wsdl",null,null,"http://mynamespace","abc");
MyInterface stub = (MyInterface)service.getStub(MyInterface.class);
stub.MyOperation(dh);
```

Attachments can also be returned from an operation, but at present only one attachment can be returned as the return parameter.

## SOAP messages with attachments - types and type mappings

By default, attachments are passed into WSIF as `DataHandlers`. If the part on the message which is the `DataHandler` maps to a `mime:part` in the WSDL, then WSIF will automatically map the fully qualified name of the WSDL type to `DataHandler.class` and set up that type mapping with the SOAP provider.

In your WSDL, you might have defined a schema for the attachment (for instance as a `binary[]`). Whether or not you have done this, WSIF silently ignores this mapping and treats the attachment as a `DataHandler`, unless you have explicitly

issued `mapType()`. WSIF lets the SOAP provider set the MIME content type based on the type of the `DataHandler`, instead of the `mime:content` type specified in the WSDL.

## Interacting with the WebSphere J2EE container

Interaction with a container is limited to the following aspects:

### Steps for this task

1. The WebSphere Application Server administrative console and WCCM allow users to define Web services to WebSphere. This is described in *Using JNDI and WSIF system management and administration*. As part of the definition of a service, the administrator may define a "Preferred Port".
2. WSIF makes log and trace calls to the WebSphere Server J2EE services, as described in *Trace and logging for WSIF*.
3. Some providers use the J2EE programming model to utilize J2EE services. The EJB provider uses JNDI and calls to remote EJBs.
4. WSIF wraps the use of container services so that when WSIF is run in an unmanaged (thin) environment, the operation can succeed.

## Running WSIF as a client

WSIF runs in the WebSphere Application Server application client container, and in similar clients from other suppliers.

To simplify the process of launching client applications in the WebSphere Application Server application client, use the `launchClient` tool as described in *(Running application clients)*.

---

## WSIF system management and administration

WSIF is provided as a stand-alone JAR file called `wsif.jar`. The JAR file contains the core WSIF classes, and the Java, EJB, SOAP over HTTP and SOAP over JMS providers. Additional providers are packaged as separate JAR files.

When you install WebSphere Application Server, `wsif.jar` is put on the WebSphere or JVM class path.

WSIF requires no further configuration. WSIF is a thin abstraction layer between application code and the relevant invocation infrastructure.

For specific information on other aspects of managing your WSIF system, see the following topics:

- Maintaining the WSIF properties file.
- Enabling security for WSIF.
- Trace and logging for WSIF.
- WSIF troubleshooting tips.
- WSIF (Web Services Invocation Framework) messages.

## Maintaining the WSIF properties file

WSIF properties are stored in a properties file (in `wsif.jar`) called `wsif.properties`. This file is kept on the class path, so that WSIF can find it and the client administrator can use it to configure WSIF.

The initial contents of `wsif.properties` are printed out below. All the possible properties are listed and described.

### Usage scenario

```
# Two properties are used to override which WSIFProvider is selected when there
# exists multiple providers supporting the same namespace URI. These properties are:
#
#   wsif.provider.default.CLASSNAME=N
#   wsif.provider.uri.M.CLASSNAME=URI
#
# CLASSNAME is the WSIFProvider class name
# N is the number of following default wsif.provider.uri.M.CLASSNAME properties
# M is a number from 1 to N to uniquely identify each wsif.provider.uri.M.CLASSNAME
# property key.
# For example the following two properties would override the default SOAP provider
# to be the Apache SOAP provider:
#
# wsif.provider.default.org.apache.wsif.providers.soap.ApacheSOAP.WSIFDynamicProvider_ApacheSOAP=1
# wsif.provider.uri.1.org.apache.wsif.providers.soap.ApacheSOAP.WSIFDynamicProvider_ApacheSOAP=\
# http://schemas.xmlsoap.org/wsdl/soap/
#

# maximum number of milliseconds to wait for a response to a synchronous request.
# Default value if not defined is to wait forever.
wsif.syncrequest.timeout=10000

# maximum number of seconds to wait for a response to an async request.
# if not defined on invalid defaults to no timeout
wsif.asyncrequest.timeout=60
```

## Enabling security for WSIF

WSIF interacts with a security manager as follows:

- WSIF runs in the current J2EE security context without modifying it.
- When WSIF is run under a J2EE container, Port implementations can utilize security context to pass on security tokens or credentials as necessary.
- WSIF implementations can automatically convert J2EE security context into appropriate context for onward services.

For WSIF to interact effectively with the WebSphere Application Server's security manager, the following permissions must be set in the `server.policy` file:

- `FilePermission` to load the WSDL (this is only required when a WSDL file is referred to using the `file:///` protocol)
- `RuntimePermission "getClassLoader"` for the current thread's context class loader (this is not required by the EJB portion).
- `RuntimePermission "accessDeclaredMembers"` (this is required by both portions for handling enterprise beans)
- `PropertyPermission` for system properties (this is required by SOAP and many others; read and write access is required for the SOAP and Java portion, only write access is required for the EJB portion)
- `NetPermission "specifyStreamHandler"` (this must be in either the SOAP and Java portion, or the EJB portion, but it need not be in both).
- `SocketPermission "host_name", "resolve"` (this is not required by the SOAP and Java portion)
- `SocketPermission "host_name:port_no", "connect"` (this is required by both portions)

where *host\_name* is your host name (for example localhost), and *port\_no* is your port number (for example 9080).

## WSIF troubleshooting tips

For information on resolving WebSphere-level problems, see (Diagnosing and fixing problems).

To identify and resolve WSIF-related problems, you can use the standard WebSphere Application Server trace and logging facilities. If you encounter a problem that you think might be related to WSIF, you can check for error messages in the WebSphere Application Server administrative console, and in the application server's `stdout.log` file. You can also enable the application server debug trace to provide a detailed exception dump.

A list of the WSIF runtime system messages, with details of what each message means, is given in Message reference for WSIF.

Here is a checklist of major WSIF activities, with advice on common problems associated with each activity:

### Create service

Handcrafted WSDL can cause numerous problems. To help ensure that your WSDL is valid, use a tool such as WebSphere Studio Application Developer (WSAD) to create your service.

### Define transport mechanism

For JMS, check that you have set up JNDI correctly, and created all the necessary queues.

For SOAP, make sure that the deployment descriptor (`dds.xml`) is correct - preferably by creating it using WSAD or similar tooling.

### Create client - Java code

Follow the correct format for creating a WSIF service, port, operation and message. For examples of correct code, see the Address Book sample.

### Compile code (client and service)

Check that the build path against code is correct, and that it contains the correct levels of JAR files.

Create a valid EAR file for your service in preparation for deployment to a web server.

### Deploy service

When you install and deploy the service EAR file, check carefully any messages given when the service is deployed.

### Server setup and start

Make sure that the WebSphere Application Server file `server.policy` (in the `/properties` directory) has the correct security settings within it. For more information see Enabling security for WSIF.

### WSIF setup

Check that file `wsif.properties` is correctly set up. For more information see Maintaining the WSIF properties file.

### Run client

Either check that you have defined the class path correctly to include



references to your client classes, WSIF JAR files and any other necessary JARs, or (preferably) run your client using the WebSphere Application Server (launchClient) tool.

Here is a list of common errors, and information on their probable causes:

**"No class definition" errors received when running client code**

This is likely to be a problem with the class path setup. Check that the relevant JAR files are included.

**"Can't find WSDL" error**

Some likely causes are:

- The application server is not running.
- The server location and port number in the WSDL are not correct.
- The WSDL is badly formed (check the error messages in the application server's `stdout.log` file).
- The application server has not been restarted since the service was installed.

You might also try the following checks:

- Can you load the WSDL into your Web browser from the location specified in the error message?
- Can you load the corresponding WSDL binding files into your Web browser?

**Your Web service's EAR file does not install correctly onto the application server.**

It is likely that the EAR file is badly formed. Verify the installation by completing the following steps:

- For an EJB binding, run the WebSphere Application Server tool `\bin\dumpnamespace`. This tool lists the current contents of the JNDI directory.
- For a SOAP over HTTP binding, open `http://pathToServer/WebServiceName/admin/list.jsp` (if you have the SOAP administration pages installed). This page lists all currently installed Web services.
- For a SOAP over JMS binding, complete the following checks:
  - Check that the Queue Manager is running.
  - Check that the necessary queues are defined.
  - Check the JNDI setup.
  - Use the `jmsadmin` tool's "display context" option to list the current JNDI definitions.
  - Check that the RPCrouter is running.

**There is a permissions problem or security error.**

Check that the WebSphere Application Server file `server.policy` (in the `/properties` directory) has the correct security settings within it. For more information see *Enabling security for WSIF*.

**Using WSIF with multiple clients causes a SOAP parsing error.**

Before you deploy a Web service to WebSphere Application Server, you must decide on the Web service's scope. The application's deployment descriptor file (`dds.xml`) includes the following line:

```
<isd:provider type="java" scope="Application" .....
```

You can set Scope to "Application" or "Session". The default setting is "Application", and this is correct if each request to the Web service does not require objects to be maintained for longer than a single instance. If Scope is set to "Application" the objects are not available to another request during the execution of the single instance, and they are released on completion. If your Web service needs objects to be maintained for multiple requests, and to be unique within each request, you must set the scope to "Session". If Scope is set to "Session", the objects are not available to another request during the life of the session, and they are released on completion of the session. If scope is set to "Application" but it should be set to "Session", you might get the following SOAP error:

```
SOAPException: SOAP-ENV:ClientParsing error, response was:
FWK005 parse may not be called while parsing.;
nested exception is:
```

```
[SOAPException: faultCode=SOAP-ENV:Client; msg=Parsing error, response was:
```

```
FWK005 parse may not be called while parsing.;
    targetException=org.xml.sax.SAXException:
FWK005 parse may not be called while parsing.]
```

## Trace and logging for WSIF

If you want to enable trace for the WSIF API within WebSphere Application Server, and have trace, stdout and stderr for the application server written to a well-known location, see (Enabling trace).

WSIF offers tracepoints at opening and closing of Ports, invocation of services, and responses from services.

To trace the WSIF API, you need to specify the following trace string:

```
wsif=all=enabled
```

WSIF also includes a SimpleLog utility through which you can run trace when using WSIF outside of WebSphere Application Server. To enable this, complete the following steps:

### Steps for this task

1. Create the file commons-logging.properties with the following contents:

```
org.apache.commons.logging.LogFactory=org.apache.commons.logging.impl.LogFactoryImpl
org.apache.commons.logging.Log=org.apache.commons.logging.impl.SimpleLog
```

2. create file simplelog.properties with the following contents:

```
org.apache.commons.logging.simplelog.defaultlog=trace
org.apache.commons.logging.simplelog.showShortLogname=true
org.apache.commons.logging.simplelog.showdatetime=true
```

3. Put both these files, and commons-logging.jar, on the class path.

### Results

The SimpleLog mechanism writes trace to file System.err.

## WSIF (Web Services Invocation Framework) messages

WebSphere system messages are logged from a variety of sources, including application server components and applications. Messages logged by application server components and associated IBM products start with a unique message identifier that indicates the component or application that issued the message. For more information about the message identifier format, see the topic (Message Format).

The rest of this topic contains a list of the WSIF runtime system messages, with details of what each message means.

**WSIF0001E: An extension registry was not found for the element type "{0}"**

**Explanation:** Parameters: {0} element type. No extension registry was found for the element type specified.

**User Response:** Add the appropriate extension registry to the port factory in your code.

**WSIF0002E: A failure occurred in loading WSDL from "{0}"**

**Explanation:** Parameters: {0} location of the WSDL file. The WSDL file could not be found at the location specified or did not parse correctly

**User Response:** Check that the location of the WSDL file is correct. Check that any network connections required are available. Check that the WSDL file contains valid WSDL.

**WSIF0003W: An error occurred finding pluggable providers: {0}**

**Explanation:** Parameters: {0} specific details about the error. There was a problem locating a WSIF pluggable provider using the J2SE 1.3 JAR file extensions to support service providers architecture. The WSIF trace file will contain the full exception details.

**User Response:** Verify that a META-INF/services/org.apache.wsif.spi.WSIFProvider file exists in a provider jar, that each class referenced in the META-INF file exists in the class path, and that each class implements org.apache.wsif.spi.WSIFProvider. The class in error will be ignored and WSIF will continue locating other pluggable providers.

**WSIF0004E: WSDL contains an operation type "{0}" which is not supported for "{1}"**

**Explanation:** Parameters: {0} name of the operation type specified. {1} name of the portType for the operation. An operation type which is not supported has been specified in the WSDL.

**User Response:** Remove any operations of the unsupported type from the WSDL. If the operation is required then make sure all messages have been correctly specified for the operation.

**WSIF0005E: An error occurred when invoking the method "{1}" . (" {0} ")**

**Explanation:** Parameters: {0} name of communication type. For example EJB or Apache SOAP. {1} name of the method that failed. An error was encountered when invoking a method on the web service using the communication shown in brackets.

**User Response:** Check that the method exists on the web service and that the correct parts have been added to the operation as described in the WSDL. Network problems might be a cause if the method is remote and so check any required connections.

**WSIF0006W: Multiple WSIFProvider found supporting the same namespace URI "{0}" . Found (" {1} ")**

**Explanation:** Parameters: {0} the namespace URI. {1} a list of the WSIFProvider found.. There are multiple org.apache.wsif.spi.WSIFProvider classes in the service provider path that support the same namespace URI.

**User Response:** A following WSIF0007I message will be issued notifying which WSIFProvider will be used. Which WSIFProvider is chosen is based on settings in the wsif.properties file, or if not defined in the properties, the last WSIFProvider found will be used. See the

wsif.properties file for more details on how to define which provider should be used to support a namespace URI.

**WSIF0007I: Using WSIFProvider "{0}" for namespaceURI "{1}"**

**Explanation:** Parameters: {0} the classname of the WSIFProvider being used. {1} the namespaceURI the provider will be used to support.. Either a previous WSIF0006W message has been issued or the SetDynamicWSIFProvider method has been used to override the provider used to support a namespaceURI.

**User Response:** None. See also WSIF0006W.

**WSIF0008W: WSIFDefaultCorrelationService removing correlator due to timeout. ID:"{0}"**

**Explanation:** Parameters: {0} the id of the correlator being removed from the correlation service. A stored correlator is being removed from the correlation service due to its timeout expiring.

**User Response:** Determine why no response has been received for the asynchronous request within the timeout period. The wsif.asyncrequest.timeout property of the wsif.properties file defines the length of the timeout period.

**WSIF0009I: Using correlation service - "{0}"**

**Explanation:** Parameters: {0} the name of the correlation service being used. This identifies the name of the correlation service that will be used to process asynchronous requests.

**User Response:** None. If a correlation service other than the default WSIF supplied one is required, ensure that it is correctly registered in the JNDI java:comp/wsif/WSIFCorrelationService namespace.

**WSIF0010E: Exception thrown while processing asynchronous response - "{0}"**

**Explanation:** Parameters: {0} the error message string of the exception. While processing the response from an executeRequestResponseAsync call an exception was thrown.

**User Response:** Use the exception error message string to determine the cause of the error. The WSIF trace will have more details on the error including the exception stack trace.

**WSIF0011I: Preferred port "{0}" was not available**

**Explanation:** Parameters: {0} the user's preferred port. The preferred port set by the user on org.apache.wsif.WSIFService is not available

**User Response:** None unless this message appears for long periods of time in which case the user might want to pick a different port as their preferred port.

---

## WSIF API

The WSIF API supports the invocation of Services defined in WSDL. WSIF is intended to be used in both WSIF clients and also in Web service intermediaries.

The WSIF API is driven by the abstract service description in WSDL; it is completely independent of the actual binding used. This makes the API more natural to work with, because it uses WSDL terms to refer to message parts, operations, and so on.

The WSIF API was designed for the WSDL usage model; to pick a port that supports the port type needed, then invoke the operation by providing the

necessary abstract input message consisting of the required parts, without worrying about how the message is mapped to a specific binding protocol.

Other Web service APIs, for example SOAP APIs, are not designed on WSDL, but for a specific binding protocol with its associated syntax; for example, target URIs and encoding styles.

The WSIF API's main interfaces are described in the following topics:

- Creating a message for sending to a port (the `WSIFMessage` interface).
- WSIF API reference: Finding a port factory or service (the `WSIFService` interface and the `WSIFServiceFactory` class).
- WSIF API reference: Using ports (the `WSIFPort` interface and the `WSIFOperation` interface).

**Note:** You must ensure that your application uses only one thread to call WSIF.

For additional technical details of the WSIF API, see the WSIF Javadoc.

## WSIF API reference: Creating a message for sending to a port

For message management (that is, message construction and parsing) the underlying API is modeled on WSDL semantics. There is a simple and direct mapping from the WSDL model to WSIF classes.

In WSDL, a Message describes the abstract type of the input or output to an operation. The corresponding WSIF class is **WSIFMessage**, which represents in memory the actual input or output of an operation. A `WSIFMessage` is a container for a set of named parts. The `WSIFMessage` interface separates the actual representation of the data from the abstract type defined by WSDL. WSDL defines Messages as XML Schema types. There are two natural ways to represent a WSDL message in a runtime environment:

- The generated Java class based on a WSDL to Java mapping such as that provided by JAX-RPC.
- The XML representation of the data, for example using SOAP Encoding.

Each option offers benefits in different scenarios. The Java class is the natural approach when WSIF is being used in a standard Java client. However, in other scenarios where WSIF is being used in an intermediary, it may be more efficient to keep a WSDL Message in the SOAP encoded format.

The style used to define messages must be consistent within the message, so all the parts in one message must be consistent. A string - `getRepresentationStyle()` - always returns `null`. This indicates that parts on this `WSIFMessage` are represented as Java objects.

Parts are added to a `WSIFMessage` with **setObjectPart** or **setTypePart**. Each part is named. Part names within a message are unique. If a part is **set** multiple times, the last time is the one that stands.

Parts are retrieved from a `WSIFMessage` by name with **getObjectPart** or **getTypePart**. If the named part does not exist, the `getXXXXPart` method returns a `WSIFException`.

Parts can be retrieved from the Message by use of Iterators through the **getParts()** and **getPartNames()** methods.

The order in which parts are set is not currently important, but the Message implementation may be more efficient if the parts are set in the parameter order specified by WSDL.

WSIFMessages are cloneable and serializable. If the parts set are not cloneable, the implementation should try to clone them using serialization. If the parts are not serializable either, then a CloneNotSupportedException will be thrown if cloning is attempted.

WSIFMessages can be sent between JVMs.

In addition to the containing parts, a WSIFMessage also has a message name. This is required for operation overloading - which is supported by WSDL and WSIF.

Here is the Javadoc for the WSIFMessage interface.

## WSIF API reference: Finding a port factory or service

To find a port you use WSIFService, which is a factory for ports. The port factory models and supports the WSDL approach in which a service is available on one or more ports. The factory hides the implementation of the port from the user. WSIF supports "dynamic" ports that are based on a particular protocol/transport and configured using the WSDL at runtime. For example, the dynamic SOAP port can invoke any SOAP service based on the WSDL description of that service. The Service allows you to hide and modify this at runtime.

Here is the WSIFService interface.

To find a service from a WSDL document at a URL, or from a code-generated codebase, you can use the WSIFServiceFactory class.

### WSIFService interface

The WSIFService is responsible for generating an instance of WSIFOperation to be used for a particular invocation of a service operation.

The WSIF service stores a list of providers that can each generate a WSIF operation for a particular WSDL binding. This service looks up providers by the provider type, so, for example, it knows about one provider that handles SOAP ports, and other providers that handle Java ports that you define. In a managed environment, the container can configure WSIFService.

Here is the Javadoc for the WSIFService interface.

A WSIFService implementation can choose a preferred port based on a number of criteria. The WSIFService implementation can set the preferred port, or it can be set by calling **setPreferredPort**.

The **getPort** method returns an instance of WSIFPort that is used to invoke a service on the port. Variants of the **getPort** method are used to define the characteristics of the port to be created. **getPort** with no arguments returns the "preferred" port. **getPort** with a string argument returns the port named by the string containing the WSDL identifier for the selected port. The return value is null if the port name is not valid.

If a port is chosen (either by WSIFService default behavior, or by setPreferredPort), then the WSIFService implementation validates that the relevant provider exists and is configured. If the provider fails this validation check, WSIFService chooses

any other port whose provider is defined. For example, if the preferred port is SOAP over JMS, but the JMS libraries are not available, WSIF chooses another port. If no preferred port is set, or the preferred port is not available, the WSIF implementation chooses the first available port listed in the WSDL.

The method `getAvailablePortNames()` returns an Iterator of Strings which is the list names of WSDL ports, filtered by the set of available providers.

The `getDefinition()` method returns the WSDL definition that was used to provide the Service. If the WSDL definition is not available, this returns a null.

### **WSIFServiceFactory class**

To find a service from a WSDL document at a URL, or from a code-generated codebase, you can use the WSIFServiceFactory class.

**Note:** When creating a WSIFService from a WSIFServiceFactory you can specify a ClassLoader object to use in locating the WSDL file. You need to do this when the WSDL file is in a JAR file. In such a case, the location of the WSDL file is specified relative to the root of the JAR file, using forward slashes with the preceding slash removed.

For example:

```
com/myCompany/wsd1/MyWSDLFile.wsd1
```

rather than

```
/com/myCompany/wsd1/MyWSDLFile.wsd1
```

Here is the Javadoc for the WSIFServiceFactory class.

The WSIFServiceFactory returns a null if no service is found with that identifier.

## **WSIF API reference: Using ports**

A WSIFPort handles the details of invoking an operation. The port provides access to the actual implementation of the service. A WSDL can provide many different WSDL Bindings, and these bindings can drive multiple Ports. The client can choose a Port, the service stub can choose a Port, or WSIF can choose a default Port.

The Port offers an interface to retrieve an Operation object. A WSIFOperation offers the ability to execute the given operation.

### **Note:**

If the Port is serialized and deserialized at a later time, then WSIF ensures that the client provides the correct information to the server to identify the instance. If the server instance is no longer available, then it is up to the server to decide whether to throw a fault or provide a new instance. That behavior can depend on the type of service.

For example, for an enterprise bean the WSIFPort stores the EJB Home, and uses it to "select" the bean before each invocation. It is the responsibility of the client to serialize or maintain the Port instance if it wants instance support. The client must create a new operation and messages for each invocation.

Here is the WSIFPort interface.

Here is the WSIFOperation interface.

### **WSIFPort interface**

The port implements a factory method for WSIFOperations.

Here is the Javadoc for the WSIFPort interface.

The `createOperation(String)` method returns a new instance of a WSIFOperation object. If the `operationName` is not valid, it will throw an exception. If the operation is overloaded, it will throw an exception. The `createOperation(String, String, String)` allows WSIF to support overloaded WSDL operations.

**Note:** You can overload based on the input parameters, but not on the output parameters.

It is the duty of the client to call the **close** method when a Port is no longer in use. In many cases, where the transport is sessionless, like HTTP, this will have no effect. However, if the Port is using a session-based protocol such as MQSeries, JMS, or ECI, this allows the Port to cache an open connection to the server and then close it as required. Responsibly-written applications will call the `close` method if appropriate.

### **WSIFOperation interface**

WSIFOperation is the runtime representation of an operation. It is responsible for invoking a service based on a particular binding. It provides methods to create input, output and fault messages, and to invoke the operation.

Here is the Javadoc for the WSIFOperation interface.

The **createInputMessage**, **createOutputMessage**, and **createFaultMessage** methods are factory methods to create the messages required by the invocation methods. All invocation methods require an input message.

The **executeRequestResponseOperation** invokes "In Out" operations.

The **executeInputOnlyOperation** invokes "In only" operations.

If the invocation method is **executeRequestResponseOperation** then an output and a fault message are instantiated and passed on the call to the **executeRequestResponseOperation** method. The output message contains the response message when the **executeRequestResponseOperation** returns true. If **executeRequestResponseOperation** returns false, then a fault occurred and is returned in the fault message.

All of these **executeNmm** methods fail with an exception if there is an error in processing the request in the WSIF provider.

The **executeRequestResponseAsync** is a feature that allows "In Out" operations to be invoked with the reply handled using an alternate thread.

Use of the **setContext** and **getContext** methods is discussed in WSIFOperation - Context.

Use of the **executeRequestResponseAsync** feature is discussed further in WSIFOperation - Asynchronous interactions.



Setting the timeouts for synchronous and asynchronous operations is discussed in *WSIFOperation - Synchronous and asynchronous timeouts*.

### **WSIFOperation - Context**

Although WSDL does not define context, a number of uses of WSIF require the ability to pass context to the Port that is invoking the service. For example, a SOAP/HTTP port may require an HTTP username and password. This information is specific to the invocation, but not a parameter of the service. In general, context is defined as a set of name-value pairs. However, because Web services tend to define the types of data using XML Schema types, WSIF represents the name-value pairs of the context using the same representation that WSIFMessages use; that is a set of named Parts, each of which equates to an instance of an XML Schema type.

The WSIFOperation methods **setContext** and **getContext** allow you to pass context information to the binding. The Port implementation can use this context - for example to update a SOAP header. There is no definition of how a Port can utilize the context.

The parameter of the **setContext** and **getContext** methods is a WSIFMessage, and this has named parts defining the context information. The WSIFConstants class defines constants for the part names that can be set in a context WSIFMessage.

The following code shows how to set the user name and password to be used for HTTP basic authentication:

```
// set a basic authentication header
WSIFMessage headers = new WSIFDefaultMessage();
headers.setObjectPart( WSIFConstants.CONTEXT_HTTP_USER, "username" );
headers.setObjectPart( WSIFConstants.CONTEXT_HTTP_PSWD, "password" );
operation.setContext( headers );
```

The WSIFOperation will ignore context parts that it does not support, for example, the above code would be ignored by the WSIF Java provider.

The WSIFConstants class includes the following constants which can be used for context part names:

- CONTEXT\_HTTP\_USER
- CONTEXT\_HTTP\_PSWD
- CONTEXT\_SOAP\_HEADERS

The HTTP header values are expected to have a type **String**, and the SOAP header value is expected to have a type of **java.util.List** which should contain entries of type **org.w3c.dom.Element**.

### **WSIFOperation - Asynchronous interactions reference**

WSIF allows asynchronous operation. In this mode of operation, the client puts the request message as part of one transaction, and carries on with the thread of execution. The response message is then handled by a different thread, with a separate transaction. The SOAP/JMS and Native JMS providers are the only WSIF providers that currently support asynchronous operation.

The WSIFPort has method **supportsAsync** to test if asynchronous operation is supported.

An asynchronous operation is initiated with the `WSIFOperation` method `executeRequestResponseAsync`. This method is designed to let an RPC method be invoked asynchronously. The method returns before the operation has completed, and the thread of execution continues.

The response to the asynchronous request is processed by one of the `WSIFOperation` methods `fireAsyncResponse` or `processAsyncResponse`.

To initiate the request, there are two forms of the `executeRequestResponseAsync` method:

```
public WSIFCorrelationId executeRequestResponseAsync
    (WSIFMessage input, WSIFResponseHandler handler)
```

and

```
public WSIFCorrelationId executeRequestResponseAsync (WSIFMessage input)
```

The first of these takes an input message and a `WSIFResponseHandler`. The `WSIFResponseHandler` is invoked on another thread when the operation completes. When using this method the client listener would call the `fireAsyncResponse` method which will call the `WSIFResponseHandler` `executeAsyncResponse` method.

Here is the Javadoc for the `WSIFResponseHandler` interface.

The other form of `executeRequestResponseAsync` that only takes a `WSIFMessage` input message does not use a response handler, and the client listener should process the response by calling the `WSIFOperation` method `processAsyncResponse`. This will update the `WSIFMessage` output and fault messages with the result of the request.

`WSIF` supports correlation between the request and response in the asynchronous request-response case. When the request is sent, the `WSIFOperation` is serialized into the `WSIFCorrelationService`. The `executeRequestResponseAsync` methods returns a `WSIFCorrelationId` object which identifies the serialized `WSIFOperation`. The client listener can use this to match a response to a particular request.

The correlation service is located with the `getCorrelationService()` method of the `WSIFCorrelationServiceLocator` class in the `org.apache.wsif.utils` package.

In a managed container a default correlation service is defined in the default JNDI namespace using the name: `java:comp/wsif/WSIFCorrelationService`. If this is not available, then `WSIF` will use the `WSIFDefaultCorrelationService`.

Here is the Javadoc for the `WSIFCorrelationService` interface.

and this is the correlator id:

```
public interface WSIFCorrelator extends Serializable {
    public String getCorrelationId();
}
```

The client must implement their own response message "listener" or MDB in order to recognize the arrival of response messages themselves. This would manage the correlation of the response message to the request and call of one of the asynchronous response processing methods. As an example of the type of thing required of a client listener, the following code fragment shows what could be in the `onMessage` method of a JMS listener:

```

public void onMessage(Message msg) {
    WSIFCorrelationService cs = WSIFCorrelationServiceLocator.getCorrelationService();
    WSIFCorrelationId cid = new JmsCorrelationId( msg.getJMSCorrelationID() );
    WSIFOperation op = cs.get( cid );
    op.fireAsyncResponse( msg );
}

```

## WSIFOperation - Synchronous and asynchronous timeouts reference

When using WSIF with JMS there are timeouts that can be set for synchronous and asynchronous operations.

Default values for these timeouts are defined in the WSIF properties file as follows:

```

# maximum number of milliseconds to wait for a response to a synchronous request.
# Default value if not defined is to wait forever.
wsif.syncrequest.timeout=10000

# maximum number of seconds to wait for a response to an async request.
# if not defined on invalid defaults to no timeout
wsif.asyncrequest.timeout=60

```

The defaults mean a synchronous request (such as a WSIFOperation `executeRequestResponseOperation` method call) will timeout after ten seconds, and an asynchronous request (such as a WSIFOperation `executeRequestResponseAsync` method call) will timeout after sixty seconds.

### Note:

The WSIFProperties class `getAsyncTimeout` method multiplies the `wsif.asyncrequest.timeout` value by 1000. This converts the value from seconds in the properties file, to milliseconds used for the timeout value.

These default values can be overridden for a given request by setting a JMS property on the operation request with the `jms:property` and `jms:propertyValue` WSDL elements. The name of the property to set should be the name of the timeout from the WSIF properties file.

For example the following sets synchronous requests to timeout after two minutes (120 seconds):

```
<jms:propertyValue name="wsif.syncrequest.timeout" type="xsd:string" value="120000"/>
```

and the following line disables asynchronous timeouts (a value of zero means wait for ever):

```
<jms:propertyValue name="wsif.asyncrequest.timeout" type="xsd:string" value="0"/>
```

**Note:** When an asynchronous timeout expires, no listener or MDB waiting for the response is notified. The asynchronous timeout is only used to tell the correlation service that the stored WSIFOperation may be deleted.

---

## WSIF: Resources for learning

Use the following links to find supplementary information about getting started with WSIF. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers

and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

- **WebSphere Version 5 Web services Handbook.** This Redbook illustrates with suitable demonstration applications how Web services can be implemented using the IBM product portfolio, especially WebSphere Application Server Version 5 and WebSphere Studio Application Developer Version 5. It includes a chapter on WSIF.
- **SOAP.** This article is a detailed overview of SOAP, which includes programming specifications.
- **Web Services Description Language.** This article is a detailed overview of Web services Description Language (WSDL), which includes programming specifications.
- **The Apache Software Foundation.** The Apache Software Foundation provides support for the Apache community of open-source software projects. Of particular interest is the **Apache Web services** project. The WSIF source code has been donated by IBM to the Apache Software Foundation, and is maintained **here** as an Apache project.
- **JSR109.** This is the Java Community Process specification for implementing Web services in Java.

See also the set of Web services links given in (services: Resources for learning).

---

## Chapter 10. IBM WebSphere UDDI Registry

Welcome to the IBM WebSphere UDDI Registry.

Use the table of contents (on the left and below) to view the various topics for a specific product or technology. Select the topic you are interested in to either open documentation locally or find information about how to locate documentation.

- "UDDI Registry terminology"
- "UDDI Registry definitions"
- "An overview of IBM UDDI Registries"
- "Migrating from the IBM WebSphere UDDI Registry on WebSphere Application Server 4.0"
- "Installing and setting up a UDDI Registry"
- "Applying Service to the UDDI Registry in a Network Deployment and single Application Server environment"
- "Reinstalling the UDDI Registry application"
- "Removing the UDDI Registry application from a deployment manager cell"
- "Removing the UDDI Registry application from a single appserver"
- "Configuring the UDDI Registry"
- "Administering the UDDI Registry"
- "UDDI user console"
- **5.0.2** "Custom Taxonomy Support in the UDDI Registry"
- "SOAP Application Programming Interface for the UDDI Registry"
- "UDDI Registry Application Programming Interface"
- "UDDI EJB Interface for the UDDI Registry"
- "UDDI troubleshooting tips"
- "Messages"
- "Running the UDDI Samples"
- "Installation Verification Program (IVP)"
- "Reporting problems with the IBM WebSphere UDDI Registry"
- "Feedback"

---

### UDDI Registry terminology

#### Syntax

When reference is made to the directory location of the WebSphere Application Server it is referred to as **<WebSphere-install-dir>** and , the directory location of the WebSphere Deployment manager as **<DeploymentManager-install-dir>**. These translate, by default, to the following locations:

#### Windows

**<WebSphere-install-dir>**

C:\Progra~1\WebSphere\AppServer\

**<DeploymentManager-install-dir>**

C:\Progra~1\WebSphere\DeploymentManager\

## 5.0+ Linux/Solaris Platforms

### 5.0.1 Linux/Solaris/HP Platforms

<WebSphere-install-dir>  
/opt/WebSphere/AppServer/

<DeploymentManager-install-dir>  
/opt/WebSphere/DeploymentManager/

#### AIX Platform

<WebSphere-install-dir>  
/usr/WebSphere/AppServer/

<DeploymentManager-install-dir>  
/usr/WebSphere/DeploymentManager/

#### z/OS Platform

<WebSphere-install-dir>  
/WebSphere390/V5R0M0/AppServer/

<DeploymentManager-install-dir>  
/WebSphere390/V5R0M0/DeploymentManager/

## UDDI Registry definitions

### Syntax

#### bindingTemplate

Technical information about a service entry point and construction specifications.

#### businessEntity

Information about the party who publishes information about a family of services.

#### businessService

Descriptive information about a particular service.

#### publisherAssertion

Information about a relationship between two parties, asserted by one or both.

#### tModel

Short for technical model.

A tModel is a data structure representing a reusable concept, such as a Web service type, a protocol used by Web services, or a category system.

tModel keys within a service description are a technical fingerprint that can be used to trace the compatibility origins of a given service. They provide a common point of reference that allows compatible services to be easily identified.

tModels are used to establish the existence of a variety of concepts and to point to their technical definitions. tModels that represent value sets such as category, identifier, and relationship systems are used to provide additional data to the UDDI core entities to facilitate discovery along a number of dimensions. This additional data is captured in keyedReferences that reside in category Bags, identifierBags, or publisherAssertions. The

tModelKey attributes in these keyedReferences refer to the value set that related to the concept or namespace being represented. The keyValues contain the actual values from that value set. In some cases keyNames are significant, such as for describing relationships and when using the general keywords value set. In all other cases, however, keyNames are used to provide a human readable version of what is in the keyValue.

---

## An overview of IBM UDDI Registries

The Universal Description, Discovery and Integration (UDDI) specification defines a way to publish and discover information about Web Services. The term 'Web service' describes specific business functionality exposed by a company, usually through an Internet connection, to allow another company, or its subsidiaries, or software program to use the service.

### Universal Business Registries (IBM UBR)

The IBM Universal Business Registry is one of a group of web-based registries that expose information about a business or other entity and its technical interfaces (or APIs). These registries are run by multiple Operator Sites, and can be used by anyone who wants to make information available about one or more businesses or entities, as well as anyone that wants to find that information. Although there are Universal Business Registries (sometimes referred to as 'public UDDI registries') hosted worldwide, including one hosted by IBM, enterprises may wish to host their own internal registries behind their firewall to better manage their internal implementation of Web Services.

For more detailed information about UDDI in general visit [ ]

### IBM WebSphere UDDI Registry

The IBM WebSphere UDDI Registry is a directory for Web Services that is implemented using the UDDI specifications. In contrast with the IBM UBR, this component of WebSphere Network Deployment is a product offering for companies or industries to implement.

A critical component of IBM's dynamic e-business infrastructure, IBM WebSphere UDDI Registry solves the problem of discovery of technical components for an enterprise and its partners by:

- Providing control, flexibility and confidentiality so that an enterprise can protect its e-business investments
- Increasing efficiency by making it easier to identify technical assets
- Leveraging existing infrastructures

For example, the IBM WebSphere UDDI Registry could be used in the following way within a large enterprise:

A company has a legacy application that provides phone numbers and Human Resources (HR) information of employees. This is turned into a Web Service and published to the registry. A developer in the same company needs to write an application for a procurement function that also needs to provide HR information to the supplier. The application should allow the supplier to have access to the employee account codes once the employee provides his name or serial number. Before Web Services, the developer had three choices:

1. Would not have known about the similar application

2. Knew about it but could not reuse due to technical barriers
3. Knew about it and reused only after significant time and negotiation

With UDDI, the developer can search for the "web service" and reuse the existing technical component in his new application for the supplier in a matter of minutes. The developer saves time and gets the application up and running sooner than he would have otherwise — increasing efficiency and saving the company time and money. The IBM WebSphere UDDI Registry is the first version 2 standard-compliant UDDI registry for private enterprise work. The IBM WebSphere UDDI Registry:

- Supports the public UDDI V2.0 standard
- Leverages the proven, reliable WebSphere Application Server technology
- Uses a relational database, such as DB2, for its persistent store.

---

## Migrating from the IBM WebSphere UDDI Registry on WebSphere Application Server 4.0

### Before you begin

If you have previously installed the IBM WebSphere UDDI Registry V1.1 (or later refreshes) on WebSphere Application Server V4.0, then you should take the following steps in order to migrate to the UDDI Registry that is available as part of WebSphere Application Server for Network Deployment V5.0.

1. If you have made any changes to the configuration properties in the file `uddi.properties`, which is located in the `properties` subdirectory of your WebSphere AppServer install tree, then you should make a copy of this file (or make a note of all the changes), so that you will be able to reapply the changes to the file after you have installed WebSphere Application Server V5.0.
2. If you have made any other configuration changes, for example to the GUI style sheets or to the SOAP interface properties, you should make a note of them, and re-apply them after upgrading to the new UDDI version.
3. Uninstall the IBM WebSphere UDDI Registry using Add/Remove Programs on Windows platforms, or `rpm -e IBMWebSphere-UDDI` on Unix Platforms. This will remove the application, but will preserve the UDDI Registry database.
4. Please note that, if you wish to continue using DB2 as the persistence store for the UDDI Registry, and if you have any data in the UDDI Registry which you wish to preserve, then you should not run the DB2 setup wizard to create the DB2 version of the UDDI Registry database, but will instead be able to continue using the database that you already have. However, if you do run the wizard, you will be prompted whether you wish to keep the database or overwrite it.
5. You can now follow the instructions on installing (or upgrading) IBM WebSphere Application Server V5 and "Installing the UDDI Registry component".
6. After completing the install procedure, you should edit the `uddi.properties` file to reflect any changes that you require to the configuration properties. You should **not** replace this file with your previous copy from the IBM WebSphere UDDI Registry V1.1 or later refreshes. You can also re-apply any other configuration changes as necessary.
7. If you are migrating from version 1.1 of the IBM WebSphere UDDI registry, then there have been a few minor changes to the EJB interface which means that you may need to modify your EJB client applications. If you are migrating



from later refreshes of the IBM WebSphere UDDI Registry (such as version 1.1.1), then you should not need to make any changes to your EJB clients. The changes for version 1.1 are when saving a new service through the EJB interface either using **saveBusiness** or **saveService**, you should not set the **serviceKey** before calling **saveBusiness** or **saveService**.

In the IBM WebSphere UDDI Registry V1.1 the term 'service type' was used to refer to a 'technical model' or 'tModel', for example, in various of the panels in the User Console. This term has now been replaced by the term 'technical model'.

---

## Installing and setting up a UDDI Registry

### Before you begin

If you wish to use the UDDI User Console using Internet Explorer as your web browser, and using SSL, then you must use Internet Explorer V5.5 with SP2 and security fix Q321232 (which must be applied in that order), or later.

If you are migrating from the IBM WebSphere UDDI Registry product that was available to run on WebSphere Application Server 4.0, then you need to read the section Migrating from the IBM WebSphere UDDI Registry on WebSphere Application Server 4.0 before installing the new product.

### Choice of database product to be used as the persistence store

The UDDI Registry application can use either DB2 or Cloudscape as the persistence store for the registry data.

- If you plan to use the UDDI Registry in production then you must use DB2 as your persistence store.
- If you plan to use the UDDI Registry for development and testing purposes, then you can choose to use Cloudscape as your persistence store. Please note that Cloudscape is not intended for production purposes.

### Steps for this task

As part of the installation of the IBM WebSphere Application Server with Network Deployment option, you are given the option to install the "UDDI Registry", which is shown under Web Services. Having selected the UDDI Registry and installed the various files that make up the application, you have two choices as to the environment in which you run it:

1. Install the UDDI Registry application into the deployment manager cell using DB2 or Cloudscape as the database in which the UDDI Registry data will be held, selecting one of the application servers within the cell in which to run the UDDI Registry.
2. Install the UDDI Registry application directly into an application server using DB2 or Cloudscape as the database in which the UDDI Registry data will be held. Please note that, if you choose this option, then the application server on which you run the UDDI Registry must not be added into a deployment manager cell, as this would cause the file synchronization within the cell to remove the application.

In most cases you will probably choose option 1, and install the UDDI Registry into a deployment manager cell, but you might find that option 2, to install the UDDI Registry into a standalone application server, is useful for development or test purposes.

**Note:**

- Several WebSphere commands are used during the following procedures, some of which must execute on the DeploymentManager and some of which must execute on the target Application server. The instructions below will distinguish which is appropriate for each command. The WebSphere commands will be found in the bin subdirectory of the appropriate WebSphere install tree. In order to ensure correct operation of these commands, you will need to do one of the following:
  - Ensure that the appropriate bin subdirectory is in your path prior to executing the command
  - Change directory to the appropriate bin subdirectory
  - Fully qualify the path to the commands
- It is recommended that you use the version of java shipped with WebSphere found in the java/bin subdirectory.

The following table lists the UDDI Registry files, and the locations into which they are placed by the install. The "Location" column shows the subdirectory under the WebSphere DeploymentManager install directory. For example, if you had installed IBM WebSphere Application Server with Network Deployment option onto a machine running Windows, and had used the default directory, then a location of "installableApps" would mean that the file had been placed into the C:\Progra~1\WebSphere\DeploymentManager\installableApps directory. For Windows platforms, read the "/" directory separator in the location column as a "\" directory separation character.

| Files                      | Purpose   | Location        |
|----------------------------|---|-----------------|
| uddi.ear                   | The UDDI Registry application itself, which is packaged and runs as an Enterprise Application   | installableApps |
| uddi.properties            | Provides configuration properties for the UDDI Registry application   | properties      |
| uddiresourcebundles.jar    | Contains system messages for the UDDI Registry application  | lib             |
| uddicloudscapeuserfunc.jar | Contains functions that are used by Cloudscape if the Cloudscape database is used with the UDDI Registry  | lib             |
| setupuddi.jacl             | Admin script to create a JDBC driver and datasource for the UDDI Registry, and to install the UDDI Registry application in a DeploymentManager Cell | UDDIReg/scripts |
| setupuddimessages.jar      | Contains setup and install messages for the UDDI Registry application   | lib             |
| removeuddi.jacl            | Admin script to undo the effects of setupuddi.jacl  | UDDIReg/scripts |
| appserverremoveuddi.jacl   | Admin script to undo the effects of appserversetupuddi.jacl   | UDDIReg/scripts |

|                           |   |                     |
|---------------------------|---|---------------------|
| appserversetupuddi.jacl   | Admin script to create a JDBC driver and datasource for the UDDI Registry, and to install the UDDI Registry application in a single, standalone, application server | UDDIReg/scripts     |
| SetupDB2UDDI.jar          | The 'UDDI DB2 Setup Wizard', to create and pre-load the UDDI Registry database if DB2 is to be used as the persistence store  | UDDIReg/scripts     |
| UDDI20 (directory)        | Cloudscape directory containing the UDDI Registry tables and pre-loaded data  | bin                 |
| uddiejbclient.jar         | Class library for use when writing an EJB client to access the UDDI Registry  | UDDIReg/ejb         |
| Various javadoc files     | JAVADOC to describe the EJB interface to the UDDI Registry  | UDDIReg/ejb/javadoc |
| UDDITaxonomyTools.jar     | Provides tools for supporting custom taxonomies with the UDDI Registry  | UDDIReg/scripts     |
| CustomTaxonomy.properties | Provides configuration properties to be used the the UDDITaxonomyTools  | UDDIReg/scripts     |

If you intend to run in a Deployment Manager Cell then complete the following task - **Installing the UDDI Registry into a deployment manager cell**

If you intend to run in a single WebSphere Application server, then complete the following task - **Installing the UDDI Registry into a single WebSphere Application Server**

What to do next

Continue with Configuring the UDDI Registry.

## **Installing the UDDI Registry into a deployment manager cell**

These instructions assume that the installation has been performed into a clean environment. If you are installing into an existing Deployment Manager cell, then steps 1 to 5 must be omitted, i.e. skip to bullet 6 immediately.

Steps for this task

1. Install the WebSphere Application Server for Network Deployment package, and select the UDDI Registry option under Web Services.
2. Install one or more base application servers which will form the cell of servers. One of these should be the application server in which you plan to run an instance of the UDDI Registry. You can run more than one instance of the UDDI Registry within a cell of servers: the UDDI Registry application name will be suffixed with the target node and server names to make it

unique within the cell (See also "Advanced use of setupuddi.jacl"), but you can only run one UDDI instance within each application server.

3. Ensure that the target application server is stopped.
4. Run startManager (startManager.sh on Unix Platforms) on the deployment manager node to start the deployment manager.

```
START dmgr_proc_name,JOBNAME=server_short_name,  
ENV=cell_short_name.node_short_name.server_short_name
```

**Note:** This command must be entered on a single line. It is split here for display purposes.

On Unix platforms you must remember to run the db2profile script before issuing the startManager.sh command. This script is located within the DB2 instance's home directory under SQLLIB and can be invoked, for example, by typing:

```
". /home/db2inst1/sqllib/db2profile"
```

5. Run addNode (addNode.sh on Unix Platforms) on each of the base application server(s) to add it as a node into the cell. (How to run addNode is described elsewhere in the InfoCenter). For example, addnode myriad - where myriad is the name of your deployment manager host.
6. Copy the uddiejbclient.jar file and the EJB javadoc directory tree from the UDDIReg/ejb subdirectory of the deployment manager install tree onto any machine(s) where you will be creating EJB clients to access the UDDI Registry.
7. If you have any global configuration properties that will be common to any UDDI Registries that you install into this cell, then you can edit the uddi.properties file in the properties subdirectory of the deployment manager install tree to set them up. (See the section on Configuring global UDDI properties for more details about the global configuration properties).
8. The UDDI Registry application is supplied with the security permissions that it requires to execute. This step explains how you can see the permissions that have been set, and change them if you feel that it is appropriate to do so. *It is recommended that you only do this if you have a thorough understanding of Java 2 security issues, and the way in which security permissions are used by WebSphere.*

The permissions for the UDDI Registry application are set within a file was.policy which is part of the uddi.ear application file. To see and change the contents of this file you should:

- a. On the deployment manager, copy the uddi.ear file from the installableApps subdirectory of the deployment manager install tree into a temporary directory.
- b. Un-jar the uddi.ear file (i.e. unpack uddi.ear using the 'jar -x' command).

For example,  
jar -x uddi.ear

(This uses the jar command in the bin subdirectory of the deployment manager, so you might need to fully qualify the path to the jar command.)

- c. You will find the was.policy file under the META-INF subdirectory that is created.

This will allow you to see the permissions which have been granted to the UDDI Registry application, and to make any changes that are necessary. Please note that if you make any errors in changing this file, then the UDDI Registry application might either fail to start, or will encounter errors when trying to execute UDDI requests.

- d. Re-jar the *uddi.ear* file using the jar command.

For example,  

```
jar -cf uddi.ear .
```

**Note:** NOTE the space and the dot after *uddi.ear*)

(This uses the jar command in the <install\_root>\java\bin subdirectory of the deployment manager, so you might need to fully qualify the path to the jar command.)

- e. Copy the new *uddi.ear* back to the *installableApps* directory.
9. *Please note that if the target application server is running, then this step will stop and restart it.* If you are planning to use Cloudscape for the database in which the UDDI Registry will be held, please read the section "Setting up the UDDI Registry to use Cloudscape within a deployment manager cell" and then return to this point. If however, you plan to use DB2, then please refer to the section "Setting up the UDDI Registry to use DB2 within a deployment manager cell" and then return to this point.
10. Ensure that the UDDI Registry is configured appropriately for your installation, as described in the section on Configuring the UDDI Registry.
11. Start, or stop and restart, the target application server. This should also start the UDDI Registry application. If not, use the admin console on the deployment manager to do so.

#### What to do next

#### Advanced use of setupuddi.jacl

A number of symbols are defined at the top of the *setupuddi.jacl* script. These allow you to control the amount of logging that is performed, and to install multiple instances of the UDDI Registry within the same cell.

The symbols which you can edit are as follows:

- **logEnabled** - default setting is 1 which causes the progress of the script to be logged. Setting this symbol to 0 causes information logging to be suppressed, with only error messages being output.
- **overwriteExisting** - default setting is 1 which causes any existing installation of the UDDI Registry application to be overwritten. Setting this symbol to 0 would cause the existing installation to be left as is, but would allow other files used by the UDDI Registry to be updated. You are recommended to only change this setting under the guidance of IBM Service.
- **appName** - default setting is *UDDIRegistry*, which will be the first part of the name used for the UDDI Registry application installed into the target server. To ensure uniqueness of application names within the cell, the full application name that will be used is <appName>.<nodeName>.<server>, where <nodeName> is the name of the target node and <server> is the name of the target server. You can choose to change the first part of this (the <appName>) portion by changing the setting of this symbol before running *setupuddi.jacl*, although it is generally recommended that you do not change this value.

Continue with Configuring the UDDI Registry.

## Setting up the UDDI Registry to use Cloudscape within a deployment manager cell

If you plan to use Cloudscape for the database in which the UDDI Registry data will be held, use this task to create and load the UDDI Registry database using Cloudscape.

### Before you begin

Please see the section "Choice of database product to be used as the persistence store" to decide which database product you should use as your persistence store before proceeding further with this step.

This task is part of the parent task and is described in Installing the UDDI Registry into a deployment manager cell. You should complete this task at the appropriate step in the parent task, then return to and complete the parent task.

If you plan to use Cloudscape for the database in which the UDDI Registry data will be held, use this task to setup and install the UDDI Registry application to use the supplied Cloudscape database.

This task configures Cloudscape on the host where you want to run the UDDI Registry. Cloudscape is supplied with WebSphere Application Server, so you should not need to install Cloudscape support.

In this task you will invoke a script called `setupuddi.jacl`, specifying the target node and application server into which the UDDI Registry is to be deployed. Please note that if the target application server is running when you invoke `setupuddi.jacl`, then the script will stop the server and will restart the server after it has completed its operations.

### Steps for this task

1. Copy the UDDI20 directory tree from the bin subdirectory of the deployment manager tree into the bin subdirectory of the target application server's install tree.
2. Create a JDBC driver and datasource to provide access to the UDDI20 Cloudscape database, and install the UDDI Registry application. This is done using the WSADMIN tool, using as input the `setupuddi.jacl` script from the UDDIReg/scripts subdirectory of the Deployment Manager. Note that this script must be run on the deployment manager node.

You should either run this script from the UDDIReg/scripts subdirectory where it is located, or copy it to some other suitable directory. Note that the WSADMIN command is located in the bin subdirectory of the deployment manager node. The syntax for calling this script for Cloudscape is:

```
wsadmin -f setupuddi.jacl
        deploymgrpath
        servername
        nodename
        discoveryURLprefix
        pathtodb
        > setupuddi.log
```

where

- *deploymgrpath* is the fully qualified pathname of the deployment manager install directory, specified using forward slashes regardless of platform; e.g. for Windows, this might be C:/Progra~1/WebSphere/deploymentManager

- *servername* is the name of the target server on which you wish to deploy the UDDI Registry, such as server1. Note the the server name entered is case sensitive.
- *nodename* is the name of the WebSphere node on which the target server runs. Note the the node name entered is case sensitive.
- *discoveryURLprefix* is the URL prefix to be used for discovery URLs. Typically this will be of the form `http://<ip-address>:9080/uddisoap/` so an example of a *discoveryURLprefix* value might be `http://mynode.mylocation.mycompany.com:9080/uddisoap/`
- *pathtodb* is the path to the UDDI20 database within the bin subdirectory of your WebSphere AppServer installation, specified using forward slashes regardless of platform; e.g. for Windows, this might be `C:/Progra~1/WebSphere/AppServer/bin/UDDI20` and for Unix platforms, it might be `/opt/WebSphere/AppServer/bin/UDDI20`
- *> setupuddi.log* is an optional parameter to direct the output to a log file as opposed to the default (which is to the screen)

For example (shown on multiple lines for publication):

```
wsadmin -f setupuddi.jacl "C:/Progra~1/WebSphere/DeploymentManager/"
server1 MYRIAD "http://myriad.headoffice.xyz.com:9080/uddisoap/"
"C:/Progra~1/WebSphere/Appserver/bin/UDDI20"
```

will install the UDDI Registry application into the server *server1* running on node *MYRIAD*, and set it up to access the Cloudscape UDDI20 database located in the bin subdirectory of the application server.

The `setupuddi.jacl` script will

- a. Create a JDBC driver named `UDDI.JDBC.Driver.<nodeName>.<server>` and a datasource named `UDDI.Datasource.<nodeName>.<server>` (where `<nodeName>` is the name of the target node and `<server>` is the name of the target server), and will replace any existing driver and datasource of that name.
- b. Check whether the UDDI Registry application is already installed and, if so, stop it and uninstall it.
- c. Update the `uddi.properties` configuration property file to configure the `discoveryURLprefix` value that you have specified and set the `persist` property as 'Cloudscape', and place this `uddi.properties` file into the location `config/cells/<currentcell>/nodes/<nodename>/servers/<servername>`. You should make any further global configuration changes using this copy of the file.
- d. Place a number of files that are needed by the UDDI Registry into the WebSphere configuration repository, and update the `ws.ext.dirs` list to reference these files.
- e. Install the UDDI Registry.

This script will deploy the UDDI Registry into the configuration under the deployment manager, and then do a `Synch` which causes it to get installed into the specified server.

### What to do next

Return to the next step in the parent task `Installing the UDDI Registry into a deployment manager cell`.

## Setting up the UDDI Registry to use DB2 within a deployment manager cell

### Before you begin

Please see the section "Choice of database product to be used as the persistence store" to decide which database product you should use as your persistence store before proceeding further with this step.

This task is part of the parent task and is described in Installing the UDDI Registry into a deployment manager cell. You should complete this task at the appropriate step in the parent task, then return to and complete the parent task.

If you plan to use DB2 for the database in which the UDDI Registry data will be held, use this task to create and load the UDDI Registry database using DB2, and to setup and install the UDDI Registry application to use the DB2 database.

This task uses the UDDI DB2 Setup Wizard to configure DB2 on the system where you want to run the UDDI Registry. Before starting this task, ensure that DB2 is installed and running on that system.

Copy the UDDIReg directory tree from the deployment manager to the target application server where DB2 will run.

The following steps should be carried out on the system on which the target application server is located (referred to below as the 'target system').

In this task you will invoke a script called setupuddi.jacl, specifying the target node and application server into which the UDDI Registry is to be deployed. Please note that if the target application server is running when you invoke setupuddi.jacl, then the script will stop the server and will restart the server after it has completed its operations.

### Steps for this task

1. On Windows, ensure that since installing DB2 you have run the usejdbc2.bat command file.

For more information about this, see "Application Building Guide" in the DB2 documentation.

2. Create and load the UDDI Registry database, called UDDI20.

**Note:** If you are migrating from an earlier version of the UDDI Registry, and your UDDI20 DB2 database already exists, then you should skip this step unless you want to replace the existing database with a new UDDI20 DB2 database. If you do choose to replace an existing database then all of your existing UDDI data will be lost. **Important:** please also note that if you do choose to replace the existing UDDI20 database, then there must not be any applications or users accessing the database at the time that you run the UDDI DB2 setup wizard.

The UDDI DB2 Setup Wizard used in this task will prompt you to provide the DB2 userid and password under which the UDDI Registry database will be created and subsequently accessed. Before starting this task, ensure that you have created an appropriate DB2 userid and password. This same userid and password must be used throughout the following steps where the DB2 userid and password is requested.



**5.0.1** If you are using a non-English installation of DB2, then please ensure you have applied PTF1 before continuing.

**5.0.1** If you wish to implement the updated SetupDB2UDDI.jacl database wizard, then also please the PTF(s) mentioned above.

**5.0.1 Note:** If you have a copy of the file SetupDB2UDDI.jar in your appserver directory, then the application of the base and Network Deployment PTFs will not update SetupDB2UDDI.jar in your appserver directory. You must apply the PTF for Network Deployment to your DeploymentManager file structure to update the SetupDB2UDDI.jar located there (in the /UDDIReg/scripts subdirectory), and then manually copy this jar to any appserver you may wish to run it on.

On Windows this should be a userid and password with administrative privileges. On Unix platforms, you should supply the userid and password of the DB2 instance in which you wish the database to be created.

To create the database you use the UDDI DB2 setup wizard, which is supplied as a jar file called SetupDB2UDDI.jar in the UDDIReg/scripts subdirectory, by following these steps:

- a. Change directory to the directory containing the file SetupDB2UDDI.jar (that is, either the UDDIReg/scripts directory in which it is supplied, or a directory on the target system into which you have copied it).
- b. In order to run the wizard, you need to first ensure that you have access from your command line to the JVM supplied with WebSphere. This is done as follows:

- On Windows, on a user with administrative privileges, in a command window type the following command:

```
was_install\bin\setupcmdline.bat
```

Where *was\_install* is the path to the directory where you installed WebSphere Application Server.

For example,

```
C:\Program Files\WebSphere\AppServer\bin\setupcmdline.bat
```

- On Unix platforms, at a command line type one of the following commands:

- If you are using bash, then as the root user run

```
./opt/WebSphere/AppServer/bin/setupCmdLine.sh
```

- If you are using csh, then as the root user run

```
source /opt/WebSphere/AppServer/bin/setupCmdLine.sh
```

where */opt/WebSphere/AppServer* is the path to the directory where you installed WebSphere Application Server.

- c. In the same command window, start the UDDI DB2 setup wizard by typing one of the following commands:

- To start a graphical user interface, type

```
java -jar SetupDB2UDDI.jar
```

- To start a text mode interface, type:

```
java -jar SetupDB2UDDI.jar -console
```

- d. Follow the prompts to work through the wizard panels.

- e. **(Optional)** If necessary, check the log files for the wizard. A log file called UDDIloadDB.log is written to the directory from which the wizard is run (but note that on Windows platforms, if you have decided not to overwrite an existing UDDI20 database, then this fact will not be logged, and the log file will not be created).
3. Create a JDBC driver and datasource to provide access to the UDDI20 DB2 database, and install the UDDI Registry application. This is done using the WASADMIN tool, using as input the setupuddi.jacl script from the UDDIReg/scripts subdirectory of the Deployment Manager. Note that this script must be run on the deployment manager node.

You should either run this script from the UDDIReg/scripts subdirectory where it is located, or copy it to some other suitable directory. Note that the WSADMIN command is located in the bin subdirectory of the deployment manager node. The syntax for this script for DB2 is:

```
wsadmin -f setupuddi.jacl
        deploymgrpath
        servername
        nodename
        discoveryURLprefix
        dbname
        db2userid
        db2password
        db2ziplocation
        > setupuddi.log
```

where

- *deploymgrpath* is the fully qualified pathname of the deployment manager install directory, specified using forward slashes regardless of platform; e.g. for Windows, this might be c:/Progra~1/WebSphere/DeploymentManager
- *servername* is the name of the target application server on which you wish to deploy the UDDI Registry, such as server1. Note that the server name is case sensitive.
- *nodename* is the name of the WebSphere node on which the target application server runs. Typically, this will be the same as the machine name. Note that the node name is case sensitive.
- *discoveryURLprefix* is the URL prefix to be used for discovery URLs. Typically this will be of the form http://<ip-address>:9080/uddisoap/ so an example of a discoveryURLprefix value might be http://mynode.mylocation.mycompany.com:9080/uddisoap/
- *dbname* is the name of the UDDI Registry database under DB2. For this parameter, you should specify UDDI20.
- *db2userid* and *db2password* are a valid DB2 userid and password with administrative privileges
- *db2ziplocation* is the path to the db2java zip file on your system, specified using forward slashes regardless of platform; e.g. for Windows, this might be C:/Progra~1/SQLLIB/java/db2java.zip
- > *setupuddi.log* is an optional parameter to direct the output to a log file as opposed to the default (which is to the screen)

For example, on Windows the command (shown here on multiple lines for publication) is:

```
wsadmin -f setupuddi.jacl "C:/Progra~1/WebSphere/deploymentManager/"
server1 MYRIAD "http://myriad.headoffice.xyz.com:9080/uddisoap/" UDDI20
db2admin secretpwd "C:/Progra~1/SQLLIB/java/db2java.zip" > setupuddi.log
```

On Unix platforms the command (shown here on multiple lines for publication) is:

```
wsadmin.sh -f setupuddi.jacl "/opt/WebSphere/deploymentManager/"
server1 MYRIAD "http://myriad.headoffice.xyz.com:9080/uddisoap/" UDDI20
db2admin secretpwd "/home/db2inst1/sql1lib/java12/db2java.zip" > setupuddi.log
```

will install the UDDI Registry application into the server `server1` running on node `MYRIAD`, and set it up to access the DB2 UDDI20 database using the userid `'db2admin'` and password `'secretpwd'`.

The `setupuddi.jacl` script will:

- a. Create a JDBC driver named `UDDI.JDBC.Driver.<nodeName>.<server>` and a datasource named `UDDI.Datasource.<nodeName>.<server>` (where `<nodeName>` is the name of the target node and `<server>` is the name of the target server), and will replace any existing driver and datasource of that name.
- b. Check whether the UDDI Registry application is already installed and, if so, stop it and uninstall it.
- c. Update the `uddi.properties` configuration file to configure the `discoveryURLprefix` value that you have specified, and to set the `persist` property as `'DB2'`, and place this `uddi.properties` file into the location `config/cells/<currentcell>/nodes/<nodename>/servers/<servername>`. You should make any further global configuration changes using this copy of the file.
- d. Place a number of files that are needed by the UDDI Registry into the WebSphere configuration repository, and update the `ws.ext.dirs` list to reference these files.
- e. Install the UDDI Registry.

#### What to do next

Return to the next step in the parent task `Installing the UDDI Registry into a deployment manager cell`.

## Installing the UDDI Registry into a single appserver

If you intend to run in a single WebSphere Application server, then complete the following task.

When you select the UDDI Registry option, then the installation will place all files that are needed to run a UDDI Registry onto the deployment manager install tree on the machine on which you are installing IBM WebSphere Application Server for Network Deployment.

To be able to run the UDDI Registry in a single application server instance in your network space you must copy these files over to the application server and then deploy the UDDI Registry. You can do this as follows:

#### Steps for this task

1. Stop the application server on which you plan to run the UDDI Registry; for example, using the command `stopServer server_name` (`stopServer.sh` on Unix Platforms)
2. Copy the `uddi.ear` file from the `installableApps` subdirectory of the deployment manager install tree into the `installableApps` subdirectory of the target application server's install tree.

3. Copy the *uddi.properties* file from the properties subdirectory of the deployment manager install tree into the properties subdirectory of the target application server's install tree.

In a subsequent step, you configure the UDDI Registry using the properties in the *uddi.properties* file.

4. Copy both the *uddiresourcebundles.jar* and the *setupuddimessages.jar* files from the lib subdirectory of the deployment manager install tree into the lib subdirectory of the target application server's install tree.
5. Optionally, if you are going to write or run code that uses the EJB interface to UDDI on another machine, then copy the *uddiejbclient.jar* file and the EJB javadoc directory tree from the UDDIReg/ejb subdirectory of the deployment manager install tree onto a location of your choice on any machines where you will be creating EJB clients to access the UDDI Registry.
6. Configure database support for the UDDI Registry database, in which the UDDI Registry will be held.

To do this, complete one of the following tasks then return this point:

- Setting up the UDDI Registry to use Cloudscape in a single AppServer
- Setting up the UDDI Registry to use DB2 in a single AppServer

**Note:** If you set up the UDDI Registry application with a JDBC driver and datasource that reference Cloudscape, but set the *persist* property in *uddi.properties* to specify DB2, **or vice versa**, then some unexpected behavior will result, such as a fatal error on deleting an entity. If this happens, you should check that the above details are not in conflict. This only applies to a UDDI Registry installation on a single appserver.

7. Ensure that the UDDI Registry is configured appropriately for your installation, as described in the section on Configuring the UDDI Registry.
8. Stop and then restart the application server.

On Unix platforms you must remember to run the *db2profile* script before issuing the *startServer.sh* command. This script is located within the DB2 instance's home directory under *SQLLIB* and can be invoked, for example, by typing:

```
". /home/db2inst1/sqllib/db2profile"
```

### What to do next

Continue with Configuring the UDDI Registry.

## Setting up the UDDI Registry to use Cloudscape in a single appserver

### Before you begin

Please see the section "Choice of database product to be used as the persistence store" to decide which database product you should use as your persistence store before proceeding further with this step.

If you plan to use Cloudscape for the database in which the UDDI Registry data will be held, use this task to setup and install the UDDI Registry application to use the supplied Cloudscape database.

This task configures Cloudscape on the host where you want to run the UDDI Registry. Before starting this task, ensure that Cloudscape is installed and running on that host, and that you can open a command window on that host.

This task, to configure Cloudscape for the UDDI Registry database. Cloudscape is supplied with WebSphere Application Server, so you should not need to install Cloudscape support.

This task, to configure Cloudscape for the UDDI Registry database, is part of the parent task to install and setup a UDDI Registry, described in Installing and Setting up a UDDI Registry. You should complete this task at the appropriate step in the parent task, then return to and complete the parent task.

To configure Cloudscape for the UDDI Registry database, complete the following steps:

#### Steps for this task

1. Copy the UDDI20 directory tree from the bin subdirectory of the deployment manager tree into the bin subdirectory of the target application server's install tree.
2. Copy the `uddicloudscapeuserfunc.jar` file from the lib subdirectory of the deployment manager install tree to the lib subdirectory of the target application server's install tree.
3. Ensure that the `persist` property in the `uddi.properties` file is set to `persist=Cloudscape`
4. Copy the `appserversetupuddi.jacl` script from the `UDDIReg/scripts` subdirectory of the deployment manager install tree to the WebSphere Application Server bin subdirectory (for example, on Windows, `C:\Progra~1\WebSphere\AppServer\bin`).
5. Change directory to the WebSphere Application Server bin subdirectory.
6. Start the application server on which the UDDI Registry is to run.  
For example, you can start the application server `server1` by typing the command:  
`startserver server1`
7. Create a JDBC driver and datasource to provide access to the UDDI20 Cloudscape database, and install the UDDI Registry application.

To do this run the `WSADMIN` tool with the script `appserversetupuddi.jacl` as input, on the target application server, using the following command syntax: (You should either run this script from the `UDDIReg/scripts` subdirectory where it is located, or copy it to some other suitable directory. Note that the `WSADMIN` command is located in the WebSphere bin subdirectory.)

```
wsadmin -f appserversetupuddi.jacl
        uddi-ear-location
        servername
        nodename
        WebSphere-lib-subdirectory
        cloudscapedbname
        > setupuddi.log
```

Where:

- `uddi-ear-location` is the fully-qualified path to the `uddi.ear` file in the `installableApps` subdirectory, specified using forward slashes regardless of platform. For example, on Windows:

```
C:/Progra~1/WebSphere/AppServer/installableApps/uddi.ear
```

- *servername* is the name of the application server on which the UDDI registry is to run; for example: server1. Note the the server name entered is case sensitive.
- *nodename* is the name of the WebSphere node on which the application server, *servername*, is running. Typically this is the machine name. Note that the node name entered is case sensitive.
- *WebSphere-lib-subdirectory* is the fully-qualified path to the WebSphere Application Server lib subdirectory, specified using forward slashes regardless of platform. For example:
  - On Windows: C:/Progra~1/WebSphere/AppServer/lib
  - On Unix: /opt/WebSphere/AppServer/lib
- *cloudscapedbname* is the fully-qualified path to the UDDI20 database within the bin subdirectory of your WebSphere AppServer installation, specified using forward slashes regardless of platform. For example:
  - On Windows, C:/Progra~1/WebSphere/AppServer/bin/UDDI20
  - On Unix, /opt/WebSphere/AppServer/bin/UDDI20
- > *setupuddi.log* is an optional parameter to direct the output to a log file as opposed to the default (which is to the screen)

The appserversetupuddi.jacl script will complete the following actions:

- a. Create a JDBC driver named UDDI.JDBC.Driver.<nodeName>.<server> and a datasource named UDDI.Datasource.<nodeName>.<server> (where <nodeName> is the name of the target node and <server> is the name of the target server), and will replace any existing driver and datasource of that name.
- b. Checks whether the WebSphere UDDI Registry application is already installed and, if so, stop the application and uninstall it.
- c. Installs the WebSphere UDDI Registry, then start it.

### What to do next

Return to the next step in the parent task Installing the UDDI Registry into a single appserver.

## **Setting up the UDDI Registry to use DB2 in a single appserver**

### Before you begin

Please see the section "Choice of database product to be used as the persistence store" to decide which database product you should use as your persistence store before proceeding further with this step.

This task, to configure DB2 for the UDDI Registry database, is part of the parent task to install and setup a UDDI Registry, described in Installing and setting up a UDDI Registry. You should complete this task at the appropriate step in the parent task, then return to and complete the parent task.

If you plan to use DB2 for the database in which the UDDI Registry data will be held, use this task to create and load the UDDI Registry database using DB2, and to setup and install the UDDI Registry application to use the database.

This task uses the UDDI DB2 setup wizard to configure DB2 on the system where you want to run the UDDI Registry. Before starting this task, ensure that DB2 is installed and running on that system.

Copy the UDDIReg directory tree from the deployment manager to the target application server where DB2 will run.

The following steps should be carried out on the system on which the target application server is located (referred to below as the 'target system').

#### Steps for this task

1. On Windows, if you are using DB2 Version 7, then ensure that since installing DB2 you have run the usejdbc2.bat command file.

For more information about this, see "Application Building Guide" in the DB2 documentation.

2. Create and load the UDDI Registry database, called UDDI20.

**Note:** If you are migrating from an earlier version of the UDDI Registry, and your UDDI20 DB2 database already exists, then you should skip this step unless you want to replace the existing database with a new UDDI20 DB2 database. If you do choose to replace an existing database then all of your existing UDDI data will be lost. **Important:** please also note that if you do choose to replace the existing UDDI20 database, then there must not be any applications or users accessing the database at the time that you run the UDDI DB2 setup wizard.

The UDDI DB2 setup wizard used in this task will prompt you to provide the DB2 userid and password under which the UDDI Registry database will be created and subsequently accessed. Before starting this task, ensure that you have created an appropriate DB2 userid and password. This same userid and password must be used throughout the following steps where the DB2 userid and password is requested.

On Windows this should be a userid and password with administrative privileges. On Unix platforms, you should supply the userid and password of the DB2 instance in which you wish the database to be created.

**5.0.1** If you are using a non-English installation of DB2, then please ensure you have applied PTF1 before continuing.

**5.0.1** Information on how to do this and where to obtain the PTF can be found here.

**5.0.1** If you wish to implement the updated SetupDB2UDDI database wizard, then also please apply the PTF(s) mentioned above.

**5.0.1** **Note:** If you have a copy of the file SetupDB2UDDI.jar in your appserver directory, then the application of the base and Network Deployment PTFs will not update SetupDB2UDDI.jar in your appserver directory. You must apply the PTF for Network Deployment to your DeploymentManager file structure to update the SetupDB2UDDI.jar located there (in the /UDDIReg/scripts subdirectory), and then manually copy this jar to any appserver you may wish to run it on.

To create the database you use the UDDI DB2 setup wizard, which is supplied as a jar file called SetupDB2UDDI.jar in the UDDIReg/scripts subdirectory, by following these steps:

- a. Change directory to the directory containing the file SetupDB2UDDI.jar (that is, either the UDDIReg/scripts directory in which it is supplied, or a directory on the target system into which you have copied it).

b. In order to run the wizard, you need to first ensure that you have access from your command line to the JVM supplied with WebSphere. This is done as follows:

- On Windows, on a user with administrative privileges, in a command window type the following command:

```
was_install\bin\setupcmdline.bat
```

Where *was\_install* is the path to the directory where you installed WebSphere Application Server.

For example, C:\Program Files\WebSphere\AppServer\bin\setupcmdline.bat

- On Unix platforms, at a command line, type one of the following commands:

- If you are using bash, then as the root user run

```
. /opt/WebSphere/AppServer/bin/setupCmdLine.sh
```

- If you are using csh, then as the root user run

```
source /opt/WebSphere/AppServer/bin/setupCmdLine.sh
```

where */opt/WebSphere/AppServer* is the path to the directory where you installed WebSphere Application Server.

c. In the same command window, start the UDDI DB2 setup wizard by issuing one of the following commands:the following command:

- To start a graphical user interface, type

```
java -jar SetupDB2UDDI.jar
```

- To start a text mode interface, type:

```
java -jar SetupDB2UDDI.jar -console
```

d. Follow the prompts to work through the wizard panels.

e. **(Optional)** If necessary, check the log files for the wizard. A log file called UDDIloadDB.log is written out into the directory from which the wizard is run (but note that, on Windows platforms, if you have decided not to overwrite an existing UDDI20 database, then this fact will not be logged, and the log file will not be created).

3. Ensure that the *persist* property in the *uddi.properties* file is set to *persist=DB2*.

4. On Unix, ensure that you have run the *db2profile* script to set up the environment for the DB2 instance that the UDDI Registry is using.

For example, type the command:

```
./home/db2inst1/sqllib/db2profile
```

5. Start the application server on which the UDDI Registry is to run.

For example, you can start the application server *server1* by issuing the command:

```
startserver server1
```

**Note:** This command must be entered on a single line. It is split here for display purposes.

6. Copy the *appserversetupuddi.jacl* script from the *UDDIReg/scripts* subdirectory of the deployment manager install tree to the WebSphere Application Server *bin* subdirectory (for example, on Windows, C:\Progra~1\WebSphere\AppServer\bin).

7. Change directory to the WebSphere Application Server *bin* subdirectory.



8. Create a JDBC driver and datasource to provide access to the UDDI20 database, and install the UDDI Registry application.

To do this run the WSADMIN tool with the script *appserversetupuddi.jacl* as input, on the target application server, using the following command syntax. (You should either run this script from the UDDIReg/scripts subdirectory where it is located, or copy it to some other suitable directory. Note that the WSADMIN command is located in the WebSphere bin subdirectory.)

```
wsadmin -f appserversetupuddi.jacl
        uddi-ear-location
        servername
        nodename
        WebSphere-lib-subdirectory
        dbname
        db2userid
        db2pwd
        db2-install-dir\location-of-db2java.zip
        > setupuddi.log
```

where

- *uddi-ear-location* is the fully-qualified path to the *uddi.ear* file in the *installableApps* subdirectory, specified **using forward slashes regardless of platform**.

For example, on Windows:

```
C:/Progra~1/WebSphere/AppServer/installableApps/uddi.ear
```

- *servername* is the name of the application server on which the UDDI registry is to run; for example: *server1*. Note that the name of the server is case sensitive.
- *nodename* is the name of the WebSphere node on which the application server, *servername*, is running. Note that the name of the node is case sensitive. Typically this will be the machine name.
- *WebSphere-lib-subdirectory* is the fully-qualified path to the WebSphere Application Server lib subdirectory, specified **using forward slashes regardless of platform**. For example:
  - On Windows: *C:/Progra~1/WebSphere/AppServer/lib*
  - On Unix or z/OS: */opt/WebSphere/AppServer/lib*
- *dbname* is the name of the UDDI Registry database under DB2. You should specify UDDI20 for this parameter
- *db2userid* and *db2pwd* are a valid DB2 userid and password with administrative privileges, as specified in an earlier step.
- *db2-install-dir* is the path under which you have installed DB2 on your system, specified **using forward slashes regardless of platform**.
  - For DB2 Version 7.2 (and later) on Windows the location is: *C:/Progra~1/SQLLIB/java12/db2java.zip* and note that the default value for *db2-install-dir* on Windows is *C:\Program Files\SQLLIB*.
  - For DB2 Version 8.1 (and later), then this parameter should be *db2-install-dir/java/db2java.zip*, and note that the default value for *db2-install-dir* on Windows is *C:\Program Files\IBM\SQLLIB*
- *> setupuddi.log* is an optional parameter to direct the output to a log file as opposed to the default (which is to the screen)

The *appserversetupuddi.jacl* will complete the following actions:

- a. Create a JDBC driver named *UDDI.JDBC.Driver.<nodeName>.<server>* and a datasource named *UDDI.Datasource.<nodeName>.<server>* (where

<nodeName> is the name of the target node and <server> is the name of the target server>, and will replace any existing driver and datasource of that name.

- b. Checks whether the WebSphere UDDI Registry application is already installed and, if so, stop the application and uninstall it.
- c. Installs the WebSphere UDDI Registry, then starts it.

#### What to do next

Return to the next step in the parent task Installing and Setting up a UDDI Registry.

---

## Reinstalling the UDDI Registry application

If you wish to reinstall the UDDI Registry then please follow the appropriate section below.

### Reinstalling into a deployment manager cell

If you wish to reinstall the UDDI Registry into the target application server, for example because you wish to alter certain aspects of its configuration using AAT, then you should rerun the setupuddi.jacl script (described in the appropriate link below

- "Setting up the UDDI Registry to use Cloudscape within a deployment cell"
- "Setting up the UDDI Registry to use DB2 within a deployment cell"

).

### Reinstalling into a single appserver

Remove the UDDI Registry application in the same manner as any other Enterprise Application and then install using the appropriate link shown below:

- Setting up the UDDI Registry to use Cloudscape in a single AppServer
- Setting up the UDDI Registry to use DB2 in a single AppServer

---

## Applying Service to the UDDI Registry in a Network Deployment and single Application Server environment

### Before you begin

When applying service to your WebSphere and UDDI Registry, you need to be aware of the following.

If you are running in a Network Deployment cell, and you apply service to the UDDI Registry (on the Deployment Manager), the updates will automatically be sent to all nodes running UDDI that are federated to that cell.

If, however, you are running in a stand-alone single Application Server (which is not known to the Deployment Manager), the updates will not automatically be sent out. In this case you must read the PTF Release Notes for instructions on how to do this.

In addition there are some updates to files and scripts that may not automatically be sent to the application servers in either a Network Deployment or single Application Server environment. Please read the PTF Release Notes for details

---

## Removing the UDDI Registry application from a deployment manager cell

### Before you begin

If you wish to completely remove the UDDI Registry application from the target application server in the deployment manager cell, then you should run the *wsadmin* (*wsadmin.sh* on Unix Platforms) script *removeuddi.jacl*, which is located in the UDDIReg/scripts directory of the deployment manager install tree.

Please note that if the target server specified on invoking *removeuddi.jacl* is running at the same time, then the script will stop the server and will restart the server after when it has completed its operations.

### Steps for this task

1.

The syntax for this script is:

```
wsadmin -f removeuddi.jacl
        servername
        nodename
        > removeuddi.log
```

Where *servername* and *nodename* are the server and node where you have deployed the UDDI Registry application. By default output will go to the screen, but, optionally, you can specify '> *removeuddi.log*' to direct output to a log file.

For example,

```
wsadmin -f removeuddi.jacl server1 myriad
```

will remove the UDDI Registry application and related files from server *server1* running in node *myriad*, and will send any messages to the screen.

---

## Removing the UDDI Registry application from a single appserver

### Before you begin

If you wish to completely remove the UDDI Registry application from a stand-alone application server then you should run the WSADMIN script *appserverremoveuddi.jacl*, which will have been installed into the UDDIReg/scripts directory when you installed the UDDI Registry as part of a Network Deployment install.

Steps for this task are:

### Steps for this task

1. The syntax for this script is:

```
wsadmin -f appserverremoveuddi.jacl
        servername
        nodename
        > removeuddi.log
```

where

- *servername* and *nodename* are the name of the stand-alone application node in which it runs (these are the names that you specified when you ran `appserversetupuddi.jacl` to install the UDDI Registry application).
- by default output will go to the screen, but, optionally, you can specify '`> removeuddi.log`' to direct the output to a log file.

For example,

```
wsadmin -f appserverremoveuddi.jacl server1 monolith
```

will remove the UDDI Registry application and related files from server `server1` running in node `monolith`, and will send any messages to the screen.

---

## Configuring the UDDI Registry

### Before you begin

The UDDI Registry is supplied as a J2EE application file, `uddi.ear`. This is installed into the WebSphere Application Server during installation. If you want to change any of its configuration properties using AAT see "Configuring SOAP properties with the AAT".

If you enable WebSphere security then to run the publish API servlet of the IBM WebSphere UDDI Registry, you also need to configure WebSphere to use HTTPS and SSL, as described in [Configuring WebSphere to use HTTPS and SSL](#)

You can configure the following aspects of the UDDI Registry:

- "Configuring global UDDI properties"
- "Modifying the database userid and password"
- "Configuring security properties"
- "Configuring the UDDI User Console (GUI) for multiple language encoding support"
- "Customizing the UDDI User Console (GUI)"
- "Configuring SOAP interface properties"
- "Configuring SOAP properties with the AAT"
- "Configuring SOAP properties in an already-deployed application"
- "Configuring WebSphere to use HTTPS and SSL"

## Configuring global UDDI properties

### Before you begin

To modify any of the global UDDI properties, you will need to edit the file called `uddi.properties`. More than one version of this file exists and the version you need to edit depends on

- whether you are in the installation phase or are updating the properties as a post installation step
- whether you are configured for a Deployment Manager or Base Application Server environment

The location of the file you should edit will be one of the following:

### Deployment Manager Configurations

1. If you are in the process of installing the UDDI Registry application for the first time into a Deployment Manager cell and wish to make some generic changes before deploying it in the cell, the *uddi.properties* file will be located in the <DeploymentManager-install-dir>/properties directory. If you are reinstalling the UDDI Registry application into a Deployment Manager cell, then you should edit the file in the location described under 2 below.

**Note:** Please note that in a Deployment Manager configuration some properties (such as *persist* and *getServletURLprefix*) are dynamically set up in the *uddi.properties* file, during subsequent installation processing.

2. If the UDDI Registry is already configured into an application server within a Deployment Manager cell (i.e. you are undertaking post installation configuration changes), then the *uddi.properties* file you should edit will be located in the configuration repository, under the Deployment Manager filing system; that is in <DeploymentManager-install-dir>config/cells/<cellname>/nodes/<nodename>/servers/<servername>, where <cellname> is the name of the Deployment Manager cell, <nodename> is the name of the node in which the application server is installed, and <servername> is the name of the application server in which you have installed the UDDI Registry.

### Application Server Configurations

1. If you are in the process of installing the UDDI Registry application into an Application Server only environment you will be advised during the installation process when to make changes to the *uddi.properties* file.  
**Note:** Please note that in contrast with the Deployment Manager configuration, UDDI properties are not dynamically set during installation processing.
2. If the UDDI Registry is already configured into a single application server which is **not** part of a Deployment Manager cell (i.e. you are undertaking post installation configuration changes), then the *uddi.properties* file will be located in the properties subdirectory of the WebSphere Application Server in which you have installed the UDDI Registry application, that is <ApplicationServer-install-dir>/properties directory.

The properties that can be changed within *uddi.properties* are as follows:

- The *dbMaxResultCount*, which is the limit on the number of rows of information that should be returned on Find requests, and will apply if the request does not specify a *maxRows* limit itself (or if it specifies a limit which exceeds this value). The initial value for this in *uddi.properties* is 100.
- The *persist*, which indicates what database is to be used as the persistence store for the UDDI Registry database. If you have installed the UDDI Registry into an application server within a Deployment Manager cell, then the *persist* property will have been set to the correct value for you. If you change this value, then you must also ensure that you have a UDDI Registry database created using the chosen database product (for more details about the UDDI Registry database, please refer to the section on "Installing the UDDI Registry"). You should also be aware that any data published to the UDDI Registry with one setting of the *persist* property will **not** be accessible when running the UDDI Registry application with a different setting for the *persist* property. The valid values for the *persist* property are:
  - *persist*=DB2  
indicating that DB2 is to be used as the persistence store
  - *persist*=Cloudscape  
indicating that Cloudscape is to be used as the persistence store

The initial value for this in *uddi.properties* is Cloudscape.

**Note:** This property is dynamically set by the *setupuddi.jacl* script when installing into a Deployment Manager cell so in this case you should not need to modify it.

- The default language to be used on a publish request as the `xml:lang` when one is not specified. The initial value for this in *uddi.properties* is `en-US`. This property must contain one of the valid `xml:lang` values.
- The UDDI site operator name. This is a string which is stored in every registry object, to indicate the operator of the UDDI Registry. The initial value for this in *uddi.properties* is `www.mycompany.com/uddi`. This property does not have any particular functional use, so its value can be set to any string that you feel is suitable.
- The maximum number of search keys that can be used on find API requests. The initial value for this in *uddi.properties* is 5.
- The `getServletURLprefix` and `getServletname` name, used to build up the discovery URL. The initial values for these are `http://localhost:9080/uddisoap/` and `get`. If you have installed the UDDI Registry into an application server within a Deployment Manager cell, then the `getServletURLPrefix` property will have been set for you using the value you specified as a parameter to the setup script. You are recommended to set suitable values for these properties before you first use the UDDI Registry.

**Note:** This property is dynamically set by the *setupuddi.jacl* script when installing into a Deployment Manager cell so in this case you should not need to modify it.

### Applying these changes to your system

In order for your changes to take effect, you must do one of the following:

- If you are in the process of installing the UDDI Registry application for the first time, you should return to the original topic and complete the installation steps. Any changes you have made will be picked up during this subsequent processing.
- If you have made post installation changes in a base Application Server only environment, you should stop and restart the UDDI Registry application using the WebSphere Administrative Console.
- If you have made post installation changes in a Deployment Manager environment you should:
  1. Execute a Full Resynchronization for the node where the UDDI Registry executes. This can be done from the WebSphere Network Deployment Administrative Console under section Systems Administration ==> Nodes. Select your node, and then click the button for 'Full Resynchronization'. It is important that you do a 'Full Resynchronization' and not just a 'Synchronize'.
  2. Stop and restart the UDDI Registry application using the WebSphere Administrative Console

## Modifying the database userid and password

### Before you begin

If you use DB2 as the persistence store for the UDDI Registry, and you need to change the database userid and/or password, you should alter the user and password values in the custom properties of the 'UDDI Datasource', which can be

edited from the WebSphere Administrative Console. The UDDI.Datasource is under datasources within the UDDI.JDBC.Driver, which is itself found under JDBC Providers under Resources. Do not alter the databaseName.

## Configuring security properties

### Before you begin

See *Configuring WebSphere to use HTTPS and SSL* for details on configuring security properties.

To run the publish API servlet of the IBM WebSphere UDDI Registry, you also need to configure WebSphere to use HTTPS and SSL, as described in *Configuring WebSphere to use HTTPS and SSL*.

## Configuring the UDDI User Console (GUI) for multiple language encoding support

### Before you begin

If you want to use multiple language encoding support in the User Console (GUI), you need to configure the application server into which the UDDI Registry application is installed with UTF-8 encoding enabled. To do this, please refer to "Configuring application servers for UTF-8 encoding" elsewhere in the WebSphere Application Server version 5 InfoCenter on enabling an application server for UTF-8.

## Customizing the UDDI User Console (GUI)

### Before you begin

The look and feel of the UDDI Console is determined by the styles defined in the *uddi\_gui.css* file which is located in the `/gui.war/theme` directory of the installed UDDI Registry application directory. The UDDI Registry application directory will be one of the following, depending on where you have installed the UDDI Registry:

- If you have installed the UDDI Registry into an application server within a Deployment Manager cell, the directory is `UDDIReg.ear` under the `/apps` directory of the configuration repository for the cell.
- If you have installed the UDDI Registry into a single application server which is not part of a Deployment Manager cell, the directory is `UDDIRgistry.ear` under the `installedApps` directory of the WebSphere Application Server in which you have installed the UDDI registry application.

The contents of this file can be edited to change the colors, fonts and font sizes according to the user's preference.

The content and layout of the UDDI User Console is provided by JSP pages, which can be customized by a programmer who is familiar with JSPs. The JSP pages are found in the `uddi.ear` enterprise application, which is under the `installedApps` subdirectory of the WebSphere AppServer installation. To locate the JSPs, you should expand the `UDDI_Registry.ear`, open the `gui.war`, and they are located under `WEB-INF` in the `pages` subdirectory. So, on a Windows system which has WebSphere installed in the default location, the JSP files will be found in "`<WebSphere-install-dir>\installedApps\UDDI_Registry.ear\gui.war\WEB-INF\pages`". These JSP pages also contain some application logic (as opposed to presentation logic) which should not be changed.

## Configuring SOAP interface properties

### Before you begin

You can configure the following SOAP interface properties

- *defaultPoolSize* - the number of SOAP parsers with which to initialize the parser pool for the SOAP interface. The can be set independently for the Publish (uddipublish) and Inquiry (uddi) APIs. For example, if you expect more inquiries than publish requests via the SOAP interface, you can set a larger pool size for the Inquiry API. The default initial size for both APIs is 10.
- The *context root* used for the Publish and Inquiry APIs, which forms a part of the URL by which they are accessed. By default this is /uddisoap.
- Whether the API is to be secure (via HTTPS) or insecure (via HTTP). The default is to use HTTPS.

To configure the following SOAP interface properties, you can use either of the following methods, as described below:

- "Configuring SOAP properties with the AAT" (the recommended option, especially for a production environment)
- "Configuring SOAP properties in an already-deployed application" for the SOAP module in the UDDI application directly. This option is faster and may be the preferred method in a test environment.

## Configuring SOAP properties with the AAT

### Before you begin

To configure SOAP properties by using the WebSphere Application Assembly tool, complete the following steps:

- Select *Update* and click on the Application icon.
- Select the *uddi.ear* file (this is placed, by the UDDI installation, into the UDDI install directory (e.g. C:\WebSphere\installableApps\uddi.ear).
- Expand the *uddi.ear* icon on the left-hand pane in the AAT.
- Expand the *Web Modules* tree.
- Expand the *uddi Soap* tree
- To change the *defaultPoolSize*, expand Web Components and then *uddipublish* (for the publish API) or *uddi* (for the inquiry API).
  - Click on *Initialization Parameters* which will show the *defaultPoolSize* parameter in the upper right-hand pane. This can be edited in the lower right-hand pane.
- To change the *context root*, click on *UDDI Soap* which will display general information about the SOAP module in the lower right-hand pane in AAT. The *context root* can be edited in this pane.
- To change the publish API to use HTTP (instead of HTTPS), click on *Security Constraints* and change the *Transport Guarantee* from Confidential to none.
- Having made any changes above, you must now save them. To do this, click on *File -> Save (or Save As)* to save your changes.
- Redeploy the *uddi.ear* to WebSphere, by first removing it and reinstalling it via the Administrator's Console.



## Configuring SOAP properties in an already-deployed application

### Before you begin

To configure SOAP properties after the UDDI application has been installed, complete the following steps:

1. The deployment descriptor for the SOAP module (*web.xml*) is found in the *WEB-INF* subdirectory of the *uddi.ear* application in the installed applications within the WebSphere install directory (for example, *<WebSphere-install-dir>\installableApps\uddi.ear\soap.war\WEB-INF*). It can be edited directly to specify the desired settings.
2. Stop and restart the WebSphere Application server for the changes to take effect.

## Configuring WebSphere to use HTTPS and SSL

### Before you begin

To support the use of secure access with the IBM WebSphere UDDI Registry, you need to configure WebSphere to use HTTPS and SSL. Please refer to the information available elsewhere in this InfoCenter for configuring SSL in WebSphere Application Server. It is assumed throughout the information for the UDDI Registry that, where SSL is used, it has been configured on port 9443.

---

## Administering the UDDI Registry

### Before you begin

This section describes the various tasks about administering the UDDI Registry.

- "Running the UDDI Registry"
- "Backing up and restoring the UDDI Registry database"

## Running the UDDI Registry

### Before you begin

The UDDI Registry is started automatically when the Application Server is started. In order to stop and restart it, use the Administrative console.

## Backing up and restoring the UDDI Registry database

### Before you begin

If you want to protect the data in your UDDI Registry database, you can backup and restore the database using the facilities of the database product. For DB2, you can do this by using the export and import functions of the DB2 Control Center. For Cloudscape you can simply use operating system tools to copy the database directory. Please refer to the database product information for more details.

The UDDI Registry database is called UDDI20, and the tables which should be backed up are:

- ADDRESS
- ADDRLINE
- BSERVICE
- BTEMPLATE

- BUSINESS
- CATEGORY
- CATEGORYBAG
- CONTACT
- DESCR
- DISCOVERYURL
- EMAIL
- EXTCATEGORY
- IDENTIFIERBAG
- INSTANCEDetail
- NAMEELEMENT
- OVERVIEWDOC
- PHONE
- PUBLISHERASSERTION
- SERVICEPROJECTION
- TMODEL
- VALIDATIONCACHE
- VALIDATIONSERVICES

---

## UDDI user console

This topic describes the layout of the UDDI user console, also referred to as the Graphical User Interface (GUI), which you can use to interact with the IBM WebSphere UDDI Registry.

For information about how to display the UDDI user console, see [Displaying the user console](#).

If you will be using the UDDI Console, then it is recommended that you configure the application server into which you have installed the UDDI Registry for UTF-8 encoding support: please refer to the section on "Configuring the UDDI User Console for multiple language encoding support".

- The user console provides a graphical user interface to the majority of the UDDI Version 2 API. It is not intended to support the full API set: there is some focus on inquiry operations, as the main purpose of the UDDI User Console is to allow users to issue inquiry requests and to familiarize themselves with general UDDI concepts. This section documents those areas for which support through the user console is not provided, together with other known restrictions to the user console.
  - General
    - Help is provided in the form of explanatory text on the screens.
    - Maximum rows cannot be specified on finds. The single maximum rows value for the registry can be set through the *dbMaxResultCount* global configuration property. For more information on setting this property see [Configuring global UDDI properties](#)
  - Find business
    - The business identifier feature is not supported.
  - Find service type
    - The business identifier feature is not supported.

- Add business
  - The business contact must be supplied as a name and role (no other information is supported).
- Add service type
  - The overview URL can be entered, but only with one description in English.
- Add service
  - There is no support for entering a Hosting Redirector, nor for adding an overviewURL.
- **Note:** The UDDI Version 2 specification states that when a tModel is deleted, it should not be physically deleted. This allows the tModel to be reinstated. One effect of this is that, if you delete a tModel using the UDDI user console, the tModel is still visible through the Show Owned Entities display.

The UDDI user console is split into three distinct areas. At the top of the screen are buttons which activate various functions in the areas below this bar. These buttons are:

**Home** returns you to the IBM WebSphere UDDI Registry welcome page

**Find** activates the Find tab on the frame below to the left

**Publish**

similarly activates the Publish tab on the frame below to the left

Below the WebSphere UDDI Registry banner the screen is split into two parts. On the left are the two tabs mentioned above, the Find and Publish tabs.

### Find tab

The Find tab is in two parts. At the top, a **Quick Find** service is provided. There are three radio buttons to enable a choice of 'service', 'business' and 'technical model' finds. Below these radio buttons is a text entry box for entering the name to search for and, beneath this, a 'Find' link to start the search. Comments are provided to show the user the wildcard character. The results of clicking on the 'Find' link are shown in the detail frame to the right.

Beneath the Quick Find is a section for **Advanced Find** functions which enables the user to choose which entity they want to perform an advanced search on. There are three links: Find services, Find businesses and Find technical models. Clicking one of these links displays the corresponding advanced search form in the frame to the right, which the user may use to enter search criteria. To initiate a Find, the user must first enter a search path (the % wildcard may be used) and then click the blue Add link to enter the search. Then click on the 'Find Services' (or 'Find Businesses/Find technical models) link below to initiate the Find operation. The **Locator** section has a link (marked in blue with the words "**Show category tree**") which displays the tree from which the user can select categories (or taxonomies). This is shown in the left-hand frame. In the advanced search form there are two links to start the search (mid-way down and at the bottom).

The results of clicking either of the two links to start the search are displayed in the same detail frame.

### Publish tab

The *Publish* link on the top banner activates the Publish tab in the navigation frame to the left. The Publish tab is split into three distinct sections.

### 1. Quick Publish Function

The top part is a **Quick Publish** section to allow the user to publish a business or technical model by name only. There are two radio buttons to enable a choice of 'business' or 'technical model'. Below these radio buttons is a text entry box for entering the name to assign to the selected entity and, beneath this, a blue 'Publish now' link to publish the entity. The results of clicking on the **Publish now** link are shown in the detail frame to the right.

### 2. Advanced Publish Functions

To publish an entity with more detail, such as with multiple names, descriptions and categories, use the **Advanced Publish** section below this. The comments below each link ('Add a business' and 'Add a technical model') describe individual functions. Clicking one of these links displays the corresponding advanced publish form in the detail frame where the user may enter details about the entity they want to publish. As in the Advanced Find functions described above, there are two links to publish a business or technical model (one towards the top of the form and the other at the bottom). Similarly the **Locator** section allows taxonomies to be shown in the left frame from which the user can select categories.

Following entry of the relevant details on the **Advanced Publish** section, the user must click on the **Publish Business** bar in order for the business to be published to the UDDI Registry.

### 3. Registered Information

Below the Advanced Publish section is a **Registered Information** section which has a link to **Show Owned Entities** in order to show the businesses, services and technical models registered to the individual user, and pending business relationships. Clicking the **Show Owned Entities** link displays the **Show Owned Entities** page in the detail frame at the right. The **Show Owned Entities** page is organized in three sections: **Registered Businesses**, **Pending Business Relationships** and **Registered Technical Models**. Each section shows the number of registered items.

#### Edit and Delete Businesses

Users can **Edit** or **Delete** businesses owned by them by clicking the appropriate links in the **Actions** column.

After an **Edit** or **Delete** function has been completed, the user **must** click on the **Update Business** bar in order to publish the changes to the UDDI Registry.

After Deleting a Business the user must confirm the deletion by clicking on the 'Delete this Business' link.

#### Adding a Service to a Business

Services are added to a business by clicking the **Add a Service** link in the **Services** column of the **Registered Businesses** section.

After the **Add a Service** function is complete, users **must** click on the **Publish Service** bar in order to publish the service to the UDDI Registry.

#### Referencing a Service from a Business

Services can also be 'referenced' by a business as if the business was the owner of the service. This 'service projection' is performed by clicking the **Reference a service** link in the **Services** column. Services associated with a business, whether they are owned or referenced, can be displayed by clicking the **Show services** link. This acts as a toggle between displaying services available for editing or deleting, and hiding them.

#### Adding a relationship to another Business

A business can be associated with another business in the UDDI Registry and this function is performed by clicking the **Add a relationship** link in the **Actions** column of the Registered Businesses section. Clicking the **Show related businesses** link in the **Actions** column displays a list of any completed business relationships.

The **Pending Business Relationships** section shows all incomplete publisher assertions, where only one party has asserted a relationship and is waiting for the other party to make the same assertion. This section reminds the user of any relationships that involve their businesses. Once both parties have asserted the same relationship between two businesses, the relationship moves from the **Pending Business Relationships** section and appears in the list of relationships displayed after clicking the **Show related businesses** link in the **Registered Businesses** section.

### Technical Models

Technical Models owned by the user are shown in the bottom **Registered Technical Models** section. As for businesses, users can Edit or Delete technical models owned by them by clicking the appropriate links in the **Actions** column.

**Note:** Users should take note that deletion of Technical Models (tModels) does **not** cause them to be physically deleted, but hidden. This is in accordance with the UDDI Registry V2.0 specifications. After deletion Technical Models are shown under the "**Shown Owned Entities**" link on the publish page but not via the Find links on the Find page. ALL other entities are deleted from the UDDI Registry in the normal way.

## Example of publishing a Business, Service and tModel with the User Console

For the example, here, we will assume a business called Mondeo Cars that sells used cars

### 1. Add the Business

Click on the Publish tab in the left hand navigation frame. Then click on 'Add a business' in the Advanced Publish in the left pane. This takes you to a 'Publish Business' pane on the right. Start by adding your Business Name in the text field labelled (Mondeo Cars in this example) and select a language and then click on the blue Add link to the right. This adds the business name (but the business is not yet published - more about which is explained later). Below the Business Name is an area called Descriptions - it allows free text to be added to describe the business - if you enter anything here you must click on the blue Add link to the right to insert the description.

The next section/area is the Locator area which can be used to describe the business according to what categories it falls into. This example uses a Used car dealership. Within the NAICS taxonomy (which you may view by clicking on 'Show category tree' and then expanding NAICS) this is a Retail Trade [44] entry which, on expansion, has Motor Vehicle and Parts Dealers [441] and, again on expansion, has automobile Dealers [4411] and Used Car Dealers [44112]. This fits the Business perfectly, so clicking on Used Car Dealers will enter the Key Name and Key Value into the business. For Checked Categorizations (such as NAICS) the Key name is not checked but the Key Value is checked. It should be noticed that for unchecked categorizations (such as 'other' or unspsc) the Key value is not checked either. If the locator field has been added, then the blue Add link must be clicked. The final area is Contacts, which can have names and role information added if required. Again, the blue Add link must be clicked after adding the relevant information.

Once all the fields are filled in to the required level, the final action is to click on the Publish Business at the bottom of the form or at the top. This causes the business to be published to the UDDI Registry and a page is displayed showing the business details.

## 2. Add a Service

From the Publish tab, there is a 'Show owned entities' link. This shows the businesses owned by the current user in the Registry and the language to be used for a particular user. For Mondeo Cars, the user will see a 'Add service' button. Clicking this button shows the Publish Service form. The top part of the form is the Service Name field. After adding this name, the user must click on the blue Add link to enter the name. As in the Adding the Business form, each subsequent part must end with the blue Add link being clicked to add that part of the information to the service. The sections are (from top to bottom, Description (a free text area), Access Points (to add link points to the Service), Locator (to add references to taxonomies to the service), and Technical Models (to associate existing tModels to the Service). After completion of those areas required, clicking on the 'Publish Service' button will Publish the service to the UDDI Registry with the current form contents.

## 3. Adding a new technical model

Clicking on the 'Add a technical model' link in the left frame opens up the Publish Technical Model form on the right. A tModel can only have one name - hence the lack of a blue Add link next to the Technical Model Name field. Beneath this field are other fields - Description (a free text area to describe the technical model), Locator (to describe the technical model with taxonomies, and an Overview URL (which gives a URL pointing to an overview document, a description of the document and a Language field). For each of these fields there is a blue Add link which must be clicked to add the relevant data. At the bottom of the form is a 'Publish Technical Model' link which will create the technical model in the UDDI Registry.

There is a Publish link at the top of the frame in each case also - after the Name section.

- "Displaying the user console"

## Displaying the user console

### Before you begin

This topic describes how to display the UDDI Registry user console (sometimes referred to as the GUI). The URL you use depends on whether or not you have enabled WebSphere security:

- If you have the WebSphere security disabled, you can access the UDDI User Console by using the following URL in your Web browser:

`http://<hostname>:9080/uddigui`

**Note:** With WebSphere security disabled, all the publish operations are performed using a userid of UNAUTHENTICATED.

- If you have WebSphere security enabled, you can access the UDDI User Console through HTTPS by using the following URL in your Web browser:

`https://<hostname>:9443/uddigui`

The User Console displays the default frameset containing the header frame, navigation frame showing find options, and details frame. When you click the link to show the publish options in the navigation frame, you are challenged for a userid and password.

If WebSphere security is enabled and you try to access a publish action via an unsecured link, e.g. clicking the publish link on the navigation frame where the User Console was opened with

`http://<hostname>:9080/uddigui`

you will be redirected to a secure logon. Inquire functions will work as expected.

## Custom Taxonomy Support in the UDDI Registry

The IBM WebSphere UDDI Registry is supplied with six published taxonomies (or categorization schemes) in the taxonomy data. Of these six, four are checked. Taxonomies can be either checked or unchecked, and this is indicated via a `keyedReference` in the `categoryBag` of the `tModel` that represents a taxonomy (a "categorization `tModel`"). These `keyedReferences` have the `tModel` key for `uddi-org:types` and are added to the `categoryBag` to further describe the behavior of the categorization `tModel`, as follows:

### checked

Marking a `tModel` with this classification asserts that it represents a categorization, identifier, or namespace `tModel` that has a validation service to check that category values are present in a specified value set.

### unchecked

Marking a `tModel` with this classification asserts that it represents a categorization, identifier, or namespace `tModel` that does not have a validation service.

In the IBM WebSphere UDDI Registry (and also in the IBM UDDI Business Registry or UBR), the validation of categories in checked taxonomies is performed against locally managed taxonomy data. Several published taxonomies are provided:

| Taxonomy name                          | Checked | Description   | tModel key   |
|--|---------|---|--|
| <code>ntis-gov:naics:1997</code>       | Yes     | Business Taxonomy: NAICS (1997 Release)   | <code>uuid:C0B9FE13-179F-413D-8A5B-5004DB8E5BB2</code> |
| <code>uddi-org:iso-ch:3166-1999</code> | Yes     | ISO 3166-1:1997 and 3166-2:1998. Codes for names of countries and their subdivisions. Part 1: Country codes. Part 2: Country subdivision codes. Update newsletters include ISO 3166-1 V-1 (1998-02-05), V-2 (1999-10-01), ISO 3166-2 I-1 (1998) | <code>uuid:4E49A8D6-D5A2-4FC2-93A0-0411D8D19E88</code> |
| <code>unspsc-org:unspsc</code>         | Yes     | Product Taxonomy: UNSPSC  | <code>uuid:CD153257-086A-4237-B336-6BDCBDCC6634</code> |
| <code>unspsc-org:unspsc:3-1</code>     | No      | Product Taxonomy: UNSPSC (Version 3.1)  | <code>uuid:DB77450D-9FA8-45D4-A7BC-04411D14E384</code> |

|                           |     |  |   |
|---------------------------|-----|--|---|
| uddi-org:types            | Yes | UDDI Type Taxonomy   | uuid:C1ACF26D-9672-4404-9D70-39B756E62AB4 |
| uddi-org:general_keywords | No  | Special taxonomy consisting of namespace identifiers and the keywords associated with the namespaces | uuid:A035A07C-F362-44DD-8F95-E2B134BF43B4 |

Taxonomy data is provided in the IBM WebSphere UDDI Registry for all the above taxonomies, apart from the general keywords taxonomy (which is unchecked). The UDDI User Console (GUI) provided with the IBM WebSphere UDDI Registry uses a shortened label for taxonomies when displayed in the taxonomy tree view, or in a pull-down list of available taxonomies as follows:

| Taxonomy Name (published) | Taxonomy name (as displayed in the GUI) |
|---------------------------|---|
| ntis-gov:naics:1997       | naics                                   |
| uddi-org:iso-ch:3166-1999 | geo                                     |
| unspsc-org:unspsc         | unspsc7                                 |
| unspsc-org:unspsc:3-1     | unspsc                                  |
| uddi-org:types            | udditype                                |
| uddi-org:general_keywords | other                                   |

This release of IBM WebSphere UDDI Registry (included with IBM WebSphere Application Server, Version 5.0.2) introduces the ability to add user-defined taxonomies, with available allowed values presented in the existing GUI taxonomy tree display. IBM WebSphere Studio Application Developer, Version 5.1 has a Web Services Explorer user interface which also allows addition and display of custom checked taxonomies. The publisher of a custom taxonomy's categorization tModel may specify a 'display name' for use in GUI implementations.

### Procedure for adding a Custom taxonomy

To add a custom taxonomy to the IBM WebSphere UDDI Registry requires you to perform two tasks: load the custom taxonomy data and publish a categorization tModel. Only when both are complete will the checked taxonomy be of practical use. Taxonomy data must be provided for validating checked taxonomies.

Taxonomy data *may* also be used by GUIs for unchecked taxonomies, but it is not a requirement and is usually only used for presentation of deprecated taxonomies, such as unspsc-org:unspsc.

If the taxonomy is checked, then any publish requests that have a categoryBag containing keyedReferences with the new categorization tModel will be validated. If there is taxonomy data corresponding to the categorization tModel in the registry database then only valid values will be accepted. If there is no taxonomy data in the database then **all** values will be rejected, and the publish request will fail. If the categorization tModel is unchecked, all values will be allowed, regardless of whether there is corresponding taxonomy data present in the UDDI Registry database.

### Suggested approach



The suggested way of introducing a new taxonomy is to:

1. Load custom taxonomy data into the UDDI Registry database using the UDDITaxonomyTools.jar utility (described below)
2. Publish the categorization tModel with a keyedReference of type 'general keywords' with keyname of 'customTaxonomy:key' and a keyValue matching the taxonomy name in the taxonomy data file (described below also)

**Note:** the SOAP and EJB interfaces will be able to make use of categorization tModels as soon as they are published. However, the UDDI Registry GUI will currently require a restart of the UDDI application because it currently gathers its list of categorizations for use in the taxonomy tree display when the application starts.

## Loading Custom Taxonomy Data

### Custom Taxonomy Data File Format

Taxonomy data is identified by a common taxonomy name, a unique code value, an optional description and a parent code which specifies its relationship with other code values. Taxonomy data must adhere to this format:

| Column name | Maximum length | Description of use  |
|-------------|----------------|---|
| name        | 8              | uniquely identifies the taxonomy within the registry  |
| code        | 32             | unique value within the taxonomy used for validation  |
| description | 128            | typically used by GUIs and optionally in the keyedReference as the keyName value              |
| parentcode  | 32             | indicates which existing code is the logical parent of this one, and is used in tree displays |

Typically columns are delimited in the taxonomy data file by '#' characters as in this example:

```
food#00#Food#00
food#10#Fruit#00
food#101#Apples#10
food#102#Oranges#10
food#103#Pears#10
food#1031#Anjou#103
food#1032#Conference#103
food#1033#Bosc#103
food#104#Pomegranates#10
food#20#Vegetables#00
food#201#Carrots#20
food#202#Potatoes#20
food#203#Peas#20
food#204#Sprouts#20
```

In the example, 'Food' is the description for the root node with child nodes of 'Fruit' and 'Vegetables' (both of these have parentcode values the same as the code value for 'Food').

The taxonomy data in the example file could then be rendered in a tree like this:

```

Food
  Fruit
    Apples
    Oranges
    Pears
      Anjou
      Conference
      Bosc
    Pomegranates
  Vegetables
    Carrots
    Potatoes
    Peas
    Sprouts

```

The file must be saved in UTF-8 format.

The following taxonomy names are reserved within the IBM WebSphere UDDI Registry and should not be used for custom taxonomy files: **naics**, **geo**, **unspsc**, **unspsc7**, **other**, **udditype**. Any attempts to publish a categorization tModel using these values for a customTaxonomy:key will be rejected. If these names are used in custom data files and the data is imported it will be indistinguishable from taxonomy data with the same name.

### UDDITaxonomyTools.jar

A utility is provided to load taxonomy data into the IBM WebSphere UDDI Registry, rename existing taxonomy data and remove existing taxonomy data, for both IBM DB2 and Cloudscape databases. The usage for each database and platform is identical:

Usage: java -jar UDDITaxonomyTools.jar {function} [options]

function:

```

-load <path>          Load taxonomy data from specified file
-rename <old> <new>  Rename existing taxonomy
-unload <name>       Unload existing taxonomy

```

options:

```

-properties <path>   Specify location of configuration file

```

**Note:** Ensure that the command window from which the UDDITaxonomyTools.jar is run is using a suitable codepage and font for displaying the characters contained in the taxonomy name.

Use of an incorrect codepage/font may result in unclear messages on a successful load, and create difficulty using the -unload and -rename options.

The following section explains in more detail how to use the utility's commands and parameters. The configuration file, if specified by the optional properties parameter, determines the database driver, authentication information and delimiters. The contents are as follows (typical data for DB2 installation shown):

| Property and example data (for DB2)   | Comments  |
|---|---|
| classpath= "c:/program files/sqllib/java12/db2java.zip; c:/tools/UDDITaxonomyTools.jar" | Classpath including database driver <i>and</i> the UDDITaxonomyTools.jar* |
| database.driver.className=com.ibm.db2.jdbc.DB2JDBC4                                     | DB2JDBC4 is the qualified classname of the database driver class          |

|  |   |
|--|---|
| database.url=c:/websphere/appserver/bin/UDDI20 | DB2 URL of the database   |
| database.userName=db2admin                     | Database userid (DB2 only)  |
| database.password=db2admin                     | Database password (DB2 only)                                      |
| column.delimiter=#                             | Column delimiter used in taxonomy data files                      |
| string.delimiter=\"                            | Field delimiter (must be different to the column.delimiter value) |

\* the classpath needs to be enclosed in quotes if the path includes space characters. Also, the UDDITaxonomyTools.jar filepath itself must be appended to the classpath (if the working directory is the same as the location of the UDDITaxonomyTools.jar then just the name is sufficient)

Filepath names should include the use of the forward-slash character (/) for all platforms.

For Cloudscape database users, the values of the following properties would be likely to be:

- classpath=c:/websphere/appserver/lib/db2j.jar; UDDITaxonomyTools.jar
- database.driver.className=com.ibm.db2j.jdbc.DB2jDriver
- database.url=jdbc:db2j:c:/ websphere/appserver/bin/uddi20

The string.delimiter is typically used where a description value contains the same character as the column delimiter character. For example, if the column.delimiter was set to ',' (comma), and there was a taxonomy description value of 'Fruits, citrus', you could include this in the taxonomy data file by setting the string.delimiter to "(double quote) and enclosing the description in quotes: 'Fruits, citrus'. Note that the quote character is escaped with a backslash to indicate the literal character is to be used.

If a properties parameter is not specified, the utility looks for and uses configuration data set in a file called customTaxonomy.properties.

**Note:** to make updates to taxonomy data in a Cloudscape database, the IBM WebSphere Application Server must be stopped to release the connection to the database.

**Note:** There is currently a limitation with UDDITaxonomyTools.jar when used with a DB2 UDDI database and multi-byte characters such as Chinese, Japanese and Korean. The maximum number of multi-byte characters is the maximum value specified earlier for name, code, description and parentcode divided by 3. For example, name can only contain values up to 8 characters in length so the maximum number of Korean characters is 2. If the taxonomy file is found to have values that exceed the limits, a message is displayed by the tool indicating the line number and column where the problem occurs. This limitation does not affect use with a Cloudscape UDDI database.

### Publishing a Checked Categorization tModel

This section describes how to publish a checked categorization tModel with the 'customTaxonomy' keyedReferences to specify which custom taxonomy data to use and a display name.

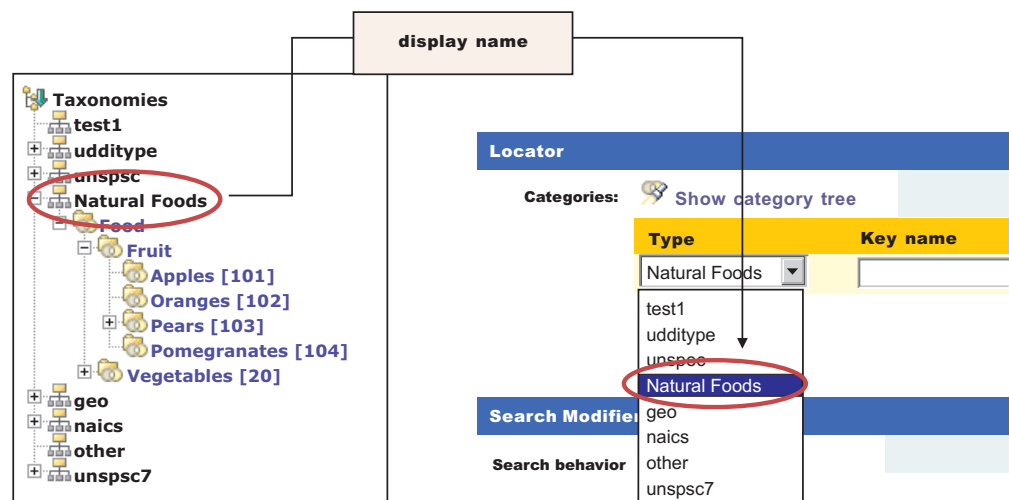
**Note:** to specify an unchecked categorization substitute the 'checked' key value with 'unchecked' or, more simply, omit the keyedReference.

Publish a tModel to the IBM WebSphere UDDI Registry with a categoryBag containing keyedReferences as follows:

| Note | tModelKey          | KeyName                                   | KeyValue                      |
|------|--------------------|---|-------------------------------|
| 1    | (uddi-org:types)   | <optional>                                | categorization                |
| 2    | (uddi-org:types)   | <optional>                                | checked                       |
| 3    | (general keywords) | urn:x-ibm:uddi:customTaxonomy:key         | <custom taxonomy displayName> |
| 4    | (general keywords) | urn:x-ibm:uddi:customTaxonomy:displayName | <custom taxonomy displayName> |

1. indicates this tModel is a categorization tModel (required)
2. indicates use of the tModel will be checked against a list of valid data (required). (Omitting this keyedReference, or explicitly specifying a value of 'unchecked' will indicate this categorization is unchecked).
3. this special use of the general keywords taxonomy, with a proprietary urn as the keyName value, defines the value used by the UDDI Registry to look up taxonomy data in its database. The value must be 1-8 (inclusive) characters long and corresponds directly with the name value in the custom taxonomy data file. Therefore, it must be unique within the registry.
4. this special use of the general keywords taxonomy, with a proprietary urn as the keyName value, defines a name for the custom taxonomy that is intended for use in GUI implementations where the full tModel name might be too long\*. The value can be 1-255 characters (inclusive) long. If this keyedReference is not supplied, the name of the tModel should be used by the GUI implementation.

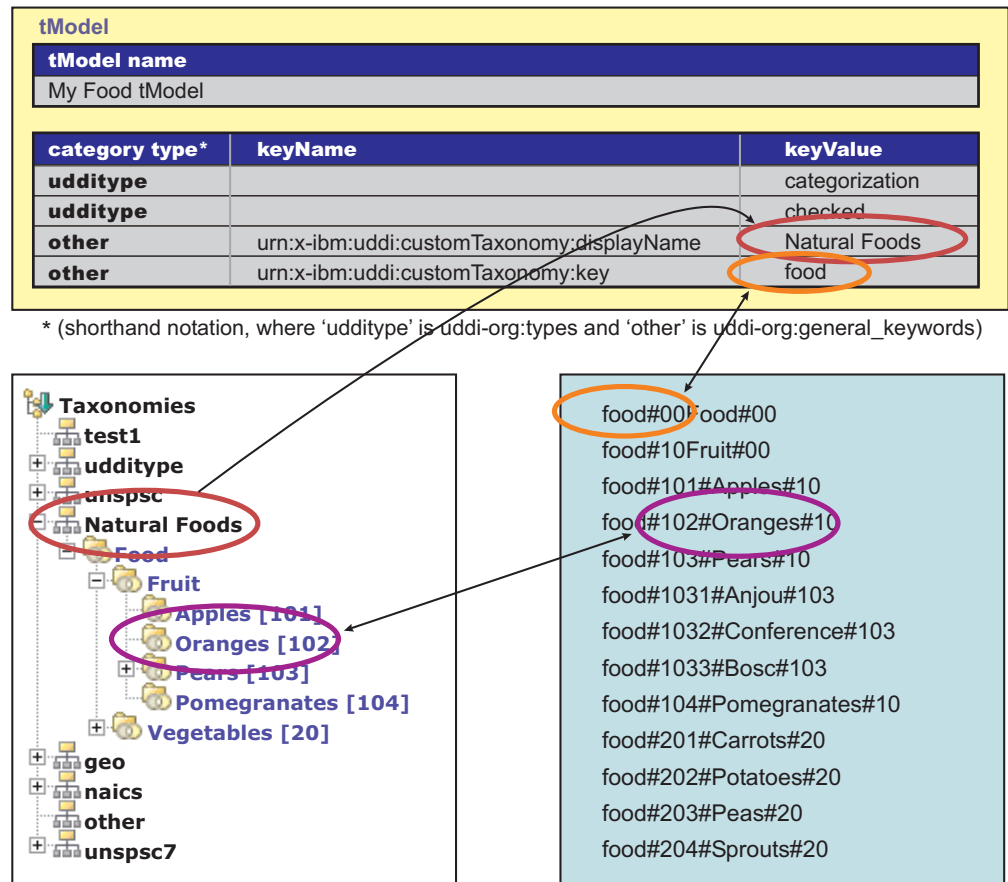
\* The displayName is intended to provide a way to label a taxonomy such that, when the UDDI GUI displays it in a taxonomy tree or in a pull-down list of available taxonomies, the meaning is clear to the user without being restricted to 8 characters and without needing to be the same as the published tModelName, which could be as long as 255 characters. An example is shown below:



Uniqueness of the urn:x-ibm:uddi:customTaxonomy:key value is validated at the time a categorization tModel is published. If it is not unique, a UDDIInvalidValueException is returned. If using a GUI to publish the tModel, an appropriate message is displayed indicating the likely cause of the problem.

The urn:x-ibm:uddi:customTaxonomy:displayName should be unique if only to avoid confusion when displayed in GUIs but this is not validated.

The relationship between the various keyedReferences, the custom taxonomy data files and use in GUIs for a categorization tModel is shown below:



As a further example, to display the label 'Delicious Victuals' in GUI displays, the categorization tModel would have a keyedReference like this:

| type  | keyName                                   | keyValue           |
|-------|---|--------------------|
| other | urn:x-ibm:uddi:customTaxonomy:displayName | Delicious Victuals |

And to link a categorization tModel to a custom taxonomy datafile with a taxonomy name of 'goodfood' the tModel's categoryBag must have a keyedReference like this:

| type  | keyName                           | keyValue |
|-------|-----------------------------------|----------|
| other | urn:x-ibm:uddi:customTaxonomy:key | goodfood |

To publish a new categorization tModel using SOAP, the message would be:

```
<save_tModel generic="2.0" xmlns="urn:uddi-org:api_v2">
  <authInfo></authInfo>>
  <tModel tModelKey="">
    <name>Natural Foods tModel</name>
    <categoryBag>
      <keyedReference tModelKey="uuid:C1ACF26D-9672-4404-9D70-39B756E62AB4"> keyValue="categorization"/>
      <keyedReference tModelKey="uuid:C1ACF26D-9672-4404-9D70-39B756E62AB4" keyValue="checked"/>
      <keyedReference tModelKey="uuid:A035A07C-F362-44DD-8F95-E2B134BF43B4"
        keyName="urn:x-ibm:uddi:customTaxonomy:key" keyValue="food"/>
      <keyedReference tModelKey="uuid:A035A07C-F362-44DD-8F95-E2B134BF43B4"
        keyName="urn:x-ibm:uddi:customTaxonomy:displayName" keyValue="Natural Foods"/>
    </categoryBag>
  </tModel>
</save_tModel>
```

**Note:** Note that 'uuid:C1ACF26D-9672-4404-9D70-39B756E62AB4' is the tModel key for uddi-org:types and 'uuid:A035A07C-F362-44DD-8F95-E2B134BF43B4' is the tModel key for uddi-org:general\_keywords.

### Validation and Error Handling

For a DB2-based IBM WebSphere UDDI Registry, custom taxonomy data can be loaded, removed and renamed using the provided utility without restarting the application (if you are using Cloudscape the application server will need to be stopped to make database updates). Removing data for which there is a corresponding checked categorization tModel will cause any use of that categorization's data to be reported as invalid.

**Note:** If an attempt is made to add data with a name that matches any of the 'internal' taxonomies, such as NAICS, GEO, etc, the request is rejected. If an attempt is made to rename or remove one of the internal taxonomies, a warning message is returned. Likewise if the user tries to rename a taxonomy to one of the reserved taxonomies, that is rejected.

The IBM WebSphere UDDI Registry user console will perform validation while a save tModel request is being built, that is, before the publish occurs. For example, if a categorization tModel with a customTaxonomy:key keyValue of 'food' already exists (in a published categorization tModel), and the user tries to add a keyedReference with the same value to the current list of keyedReferences, the following message is displayed:

```
Advice: The 'urn:x-ibm:uddi:customTaxonomy:key' value of 'food' is already in use by
another categorization tModel. Enter a unique value
```

Similarly, only one of each of the customTaxonomy:key and customTaxonomy:displayName keyedReferences are allowed. For example, if the user tries to add two customTaxonomy:displayName keyedReferences they will get the message:

```
Advice: Only one 'urn:x-ibm:uddi:customTaxonomy:displayName' key name is allowed
for the 'Other' taxonomy
```

If the customTaxonomy:key keyedReference is valid and unique at the time it is added to the save\_tModel request, the keyedReference is further validated when the user makes the publish request, to ensure that another session has not successfully published a categorization tModel with the same customTaxonomy:key. In this case, the user is returned to the Publish Technical Model page.

If a keyedReference containing a keyName value that starts with 'urn:x-ibm:uddi:customTaxonomy:' is followed by anything other than 'key' or 'displayName', the following message is displayed:

Advice: Only key name values of 'urn:x-ibm:uddi:customTaxonomy:displayName' and 'urn:x-ibm:uddi:customTaxonomy:key' are supported.

For SOAP, UDDI4J, and EJB initiated requests where the save\_tModel message may have multiple tModels, if any one of the tModels is a categorization tModel and it fails validation, the request fails with a UDDIInvalidValueException (plus additional information explaining the likely cause), and none of the tModels is published. For example, if a publish request includes a customTaxonomy:key keyedReference with a keyValue that matches the customTaxonomy:key keyValue of an existing categorization tModel, the following UDDIInvalidValueException is thrown, with message:

```
E_invalidValue (20200) A value that was passed in a keyValue attribute did not pass validation. This applies to checked categorizations, identifiers and other validated code lists. The error text will clearly indicate the key and value combination that failed validation. Invalid 'customTaxonomy:dbKey' keyValue [naics] in keyedReference. KeyValue already in use by tModelKey[UUID:C0B9FE13-179F-413D-8A5B-5004DB8E5BB2]
```

The customTaxonomy:key and customTaxonomy:displayName keyValue values are validated. For example, a publish categorization tModel request with a keyedReference including a customTaxonomy:key of 'toolongdbkey' was attempted, the following UDDIInvalidValueException is thrown, with message:

```
E_invalidValue (20200) A value that was passed in a keyValue attribute did not pass validation. This applies to checked categorizations, identifiers and other validated code lists. The error text will clearly indicate the key and value combination that failed validation. Invalid 'customTaxonomy:key' keyValue [toolongdbkey] in keyedReference. tModelKey[]
```

If a categorization tModel is edited in the user console, or republished via SOAP, UDDI4J or EJB, such that it is no longer a categorization tModel (ie the categorization keyedReference is removed), then that tModel is removed from the internal store of categorization tModels, and its customTaxonomy:key value, if present, is available for use by new categorization tModels.

---

## SOAP Application Programming Interface for the UDDI Registry

Access to the SOAP API will by default be available at:

[ ] or

[ ]

Where 'localhost' is the address by which your WebSphere server is known. If security is enabled on your WebSphere server, the publishapi will also be protected by basic-authentication. By default, when security is enabled, the publishapi is restricted to HTTPS, this is to ensure the confidentiality and security of your data whilst in transit to UDDI. If you do not wish to use SSL, when security is enabled, you will have to modify the jar file using AAT to remove the CONFIDENTIAL restriction placed upon the publish URLs. For more information about this topic, see the section on **Configuring SOAP properties with the AAT** If you normally access your WebSphere server via a web server, you will need to ensure the plugin configuration for the WebSphere plugin on the web server has been updated since installing UDDI. This will then allow access to the UDDI SOAP API through the URLs

[ ]or

[ ]

Where 'localhost' is the address by which your web server is accessed. Note that if you plan on accessing UDDI via a web server in this manner, that the samples will require modification to inform them of the SSL certificates used by your web server, so that the samples can make SSL connections to the web server. It is beyond the scope of this document to cover the many variants available on web server/WebSphere/java SSL configurations

- "Programming the UDDI SOAP API"
- "SOAP API error handling tips in the UDDI Registry"

## Programming the UDDI SOAP API

To use the SOAP API construct a properly formed UDDI message within the body of a SOAP request, and send it using HTTP POST to the URL of the API which the request relates to. The response will be returned within the body of the HTTP reply. Although the samples are written in Java, you may use other programming languages to create your SOAP client, providing you still send requests compliant to the SOAP specification. Valid UDDI requests should conform to the UDDI schema, and be as detailed within the UDDI standard documentation available from:

[ ]

For more information on using the SOAP API, refer to [ ] section within this InfoCenter.

## SOAP API error handling tips in the UDDI Registry

When using the SOAP API there are three main categories that may cause an error to be returned:-

- An invalid/incorrect request being sent to the SOAP API. e.g. Incorrectly formed XML, Badly formed UDDI requests, Non-schema compliant requests.
- Invalid business logic within a SOAP API request. e.g. Attempting to delete a business that does not exist.
- Problems occurring while processing a valid request. e.g. Server connection to database failure.

In each of these cases, an error will be returned to the client that made the request, which will attempt to explain further what the problem was.

---

## UDDI Registry Application Programming Interface

The IBM WebSphere UDDI Registry fully supports the Application Programming Interface (API) specification which can be viewed by visiting [ ]. Any changes from this specification are documented within the IBM WebSphere UDDI Registry information.

- The Inquiry API
- The Publish API

## Inquiry API for the UDDI Registry

The Inquiry API provides three forms of query that follow broadly used conventions which match the needs of software traditionally used within registries.



- "Browse pattern for the UDDI Registry"
- "Drilldown pattern for the UDDI Registry"
- "Invocation pattern for the UDDI Registry"
- "Inquiry API functions in the UDDI Registry"

## Browse pattern for the UDDI Registry

Software that allows people to explore and examine data - especially hierarchical data - requires browse capabilities. The browse pattern characteristically involves starting with some broad information, performing a search, finding general result sets and then selecting more specific information for drill-down.

The UDDI API specifications accommodate the browse pattern by way of the *find\_xx* API calls. These calls form the search capabilities provided by the API and are matched with summary return messages that return overview information about the registered information that is associated with the inquiry message type and the search criteria specified in the inquiry.

A typical browse sequence might involve finding whether a particular business you know about has any information registered. This sequence would start with a call to *find\_business*, perhaps passing the first few characters of a business name that you already know. This returns a *businessList* result. This result is overview information (keys, names and descriptions) derived from the registered *businessEntity* information, matching on the name fragment that you provided. If you spot the business you are looking for within this list, you can drill down into the corresponding *businessService* information, looking for particular technical models (e.g. purchasing, shipping, etc) using the *find\_service* API call. Similarly, if you know the technical *fingerprint* (tModel signature) of a particular software interface and want to see if the business you've chosen provides a web service that supports that interface, you can use the *find\_binding* inquiry message.

## Drilldown pattern for the UDDI Registry

When you have a key for one of the four main data types managed by a UDDI registry, you can use that key to access the full registered details for a specific data instance. The UDDI data types are *businessEntity*, *businessService*, *bindingTemplate* and *tModel*. You can access the full registered information for any of these structures by passing a relevant key type to one of the *get\_xx* API calls.

Continuing the example from the previous section on browsing, one of the data items returned by all of the *find\_x* return sets is key information. In the case of the business we were interested in, the *businessKey* value returned within the contents of a *businessList* structure can be passed as an argument to *get\_businessDetail*. The successful return to this message is a *businessDetail* message containing the full registered information for the entity whose key value was passed. This will be a full *businessEntity* structure.

## Invocation pattern for the UDDI Registry

In order to prepare an application to take advantage of a remote Web service that is registered within the UDDI registry by other businesses or entities, you need to prepare that application to use the information found in the registry for the specific service being invoked.

The *bindingTemplate* data obtained from the UDDI registry represents the specific details about an instance of a given interface type, including the location at which a program starts interacting with the service. The calling application or program

should cache this information and use it to contact the service at the registered address whenever the calling application needs to communicate with the service instance. In previously popular remote procedure technologies tools have automated the tasks associated with caching (or hard coding) location information. Problems arise however when a remote service is moved without any knowledge on the part of the callers. Moves occur for a variety of reasons, including server upgrades, disaster recovery, and service acquisition and business name changes.

When a call fails using cached information previously obtained from a UDDI Registry, the proper behavior is to query the UDDI Registry for fresh bindingTemplate information. The proper call get\_bindingDetail passing the original bindingKey value. If the data returned is different from the cached information, the service invocation should automatically retry the invocation using the fresh information. If the result of this retry is successful, the new information should replace the cached information.

By using this pattern with Web services, a business using a UDDI Registry can automate the recovery of a large number of partners without undue communication and coordination costs. For example, if a business has activated a disaster recovery site, most of the calls from partners will fail when they try to invoke services at the failed site. By updating the UDDI information with the new address for the service, partners who use the invocation pattern will automatically locate the new service information and recover without further administrative action.

## Inquiry API functions in the UDDI Registry

These messages represent inquiries that can be made of the UDDI Registry. These messages all behave synchronously.

The queries available are:

### **find\_binding**

Used to locate specific bindings within a registered businessService. Returns a bindingDetail message.

### **find\_business**

Used to locate information about one or more businesses. Returns a businessList message.

### **find\_relatedBusinesses**

Used to locate information about businessEntity registrations that are related to a specific business entity whose key is passed in the inquiry. The Related Businesses feature is used to manage registration of business units and subsequently relate them based on organizational hierarchies or business partner relationships. Returns a relatedBusinessList message.

### **find\_service**

Used to locate specific services within a registered businessEntity. Returns a serviceList message.

### **find\_tModel**

Used to locate one or more tModel information structures. Returns a tModelList structure.

### **get\_bindingDetail**

Used to get full bindingTemplate information suitable for making one or more service requests. Returns a bindingDetail message.

**get\_businessDetail**

Used to get the full businessEntity information for one or more businesses or organizations. Returns a businessDetail message.

**get\_serviceDetail**

Used to get full details for a given set of registered businessService data. Returns a serviceDetail message.

**get\_tModelDetail**

Used to get full details for a given set of registered tModel data. Returns a tModelDetail message.

## Accessible query values in the UDDI Registry

A list of the accessible queries for the UDDI Registry is given here.

Accessible queries within the UDDI Registry are:

**find\_binding**

used to locate specific bindings within a registered businessService. Returns a bindingDetail message that contains zero or more bindingTemplate structures matching the criteria specified in the argument list.

**find\_business**

used to locate information about one or more businesses. Returns a businessList message that matches the conditions specified in the arguments.

**find\_relatedBusinesses**

used to locate information about businessEntity registrations that are related to a specific business entity whose key is passed in the inquiry. The Related Businesses feature is used to manage registration of business units and subsequently relate them based on organizational hierarchies or business partner relationships. Returns a relatedBusinessList message containing results that match the conditions specified in the arguments.

**find\_Service**

used to locate specific services within a registered businessEntity. Returns a serviceList message that matches the conditions specified in the arguments.

**find\_tModel**

used to locate a list of tModels that match a set of specified criteria. The response will be a list of abbreviated information about registered tModel data that matches the criteria specified. The result will be returned in a tModelList message.

**get\_bindingDetail**

used to requesting the run-time bindingTemplate information for the purpose of invoking a registered business API. Returns a bindingDetail message.

**get\_businessDetail**

used to return complete businessEntity information for one or more specified businessEntity registrations matching on the businessKey values specified.

**get\_businessDetailExt**

used to return extended businessEntity information for one or more specified businessEntity registrations. This message returns exactly the

same information as the `get_businessDetail` message, but may contain additional attributes if the source is an external registry that is compatible with the API specification.

**get\_serviceDetail**

used to request full information about a known `businessService` structure. Returns a `serviceDetail` message.

**get\_tModelDetail**

`get_tModelDetail`

For full details of the syntax of the above queries, please refer to the API Specification at [ ]

## Publish API for the UDDI Registry

The messages in this section represent commands that are used to publish and update information contained in a UDDI registry. The messages defined in this section all behave synchronously.

The Publishing API calls defined that UDDI operators support are:

**add\_publisherAssertions**

this call causes one or more `publisherAssertion` to be added to an individual publisher's assertion collection.

**delete\_binding**

causes one or more instances of `bindingTemplate` data to be deleted from the UDDI registry.

**delete\_business**

used to remove one or more business registrations and all direct contents from a UDDI registry.

**delete\_publisherAssertions**

causes one or more `publisherAssertion` elements to be removed from a publisher's assertion collection.

**delete\_service**

is used to remove one or more `businessService` elements from the UDDI registry and from its containing `businessEntity` parent.

**delete\_tModel**

is used to logically delete one or more `tModel` structures. Logical deletion hides the deleted `tModels` from `find_tModel` result sets but does not physically delete it, so it is returned on a `get_registeredInfo` request.

**discard\_authToken**

is used to inform an Operator Site that the authentication token is to be discarded, effectively ending the session. Subsequent calls that use the same `authToken` will be rejected. This message is optional for Operator Sites that do not manage session state or that do not support the `get_authToken` message.

**get\_assertionStatusReport**

this call provides administrative support for determining the status of current and outstanding publisher assertions that involve any of the business registrations managed by the individual publisher account. Using this message, a publisher can see the status of assertions that they have made, as well as see assertions that others have made that involve `businessEntity` structures controlled by the calling publisher account.

**get\_authToken**

the call used to obtain an authentication token. Authentication tokens are opaque values that are required for all other publisher API calls. This message is not required for Operator Sites that have an external mechanism defined for users to get an authentication token. This API is provided for implementations that do not have some other method of obtaining an authentication token or certificate, or that choose to use user ID and password based authentication.

**get\_publisherAssertions**

this is used to obtain the full set of publisher assertions that are associated with an individual publisher account. Publisher assertions are used to control publicly visible business relationships.

**get\_registeredInfo**

this call is used to get an abbreviated list of all businessEntity and tModel data that are controlled by the individual associated with the credentials passed.

**save\_binding**

is used to save or update a complete bindingTemplate element. this message can be used to add or update one or more bindingTemplate elements as well as the container/contained relationship that each bindingTemplate has with one or more existing businessService elements.

**save\_business**

this is used to save or update information about a complete businessEntity element. This API has the broadest scope of all the save\_xx API calls in the publisher API, and can be used to make sweeping changes to the published information for one or more businessEntity elements controlled by an individual.

**save\_service**

the call used to add or update one or more businessService elements exposed by a specified businessEntity.

**save\_tModel**

this call adds or updates one or more registered tModel elements.

**set\_publisherAssertions**

this call is used to manage all of the tracked relationship assertions associated with an individual publisher account.

For full details of the syntax of the above queries, please refer to the API Specification at [ ].

---

## UDDI EJB Interface for the UDDI Registry

This section describes how to use the EJB application programming interface (API) of the IBM WebSphere UDDI Registry component to publish, find and delete UDDI entries.

The necessary client classes are contained in the *uddiejbclient.jar* file in the *ejb* subdirectory of the *UDDIReg* directory under the WebSphere appserver directory tree.

The javadoc for the EJB API is contained in the javadoc directory tree under the *ejb* subdirectory of the *UDDIReg* directory under the WebSphere appserver directory tree.

The EJB API is contained in two stateless session beans, one for the Inquiry API (`com.ibm.uddi.ejb.InquiryBean`) and one for the Publish API (`com.ibm.uddi.ejb.PublishBean`), whose public methods form an EJB interface for the UDDI Registry. All the public methods on the `InquiryBean` correspond to UDDI Inquiry API functions, and all the public methods on the `PublishBean` correspond to UDDI Publish API functions. (Not all UDDI API functions are implemented, e.g. `get_authToken`, `discard_authToken`, `get_businessDetailExt`, etc.) For version 1 of the UDDI registry, the EJB component supports only UDDI v2.0.

The two EJBs use container-managed transactions. The transaction attribute for the methods of the `InquiryBean` is `NotSupported`, and for the methods of the `PublishBean` it is `Required`. You should not change the transaction attributes as this could result in undesirable behavior.

Within each interface there are groups of overloaded methods that correspond to the operations in the UDDI 2.0 specification. There is a separate method for each major variation in function. For example, the single UDDI 2.0 operation `find_business` is represented by 10 variations of `findBusiness` methods, with different variations for finding by name, finding by categoryBag etc.

The arguments for the EJB interface methods are java objects in the package `com.ibm.uddi.datatypes`. Roughly speaking, there is a one-one correspondence between classes in this package and elements of the UDDI v2.0 XML schema. Exceptions to this are, for example, where UDDI XML elements can be represented by a single String. (See Package `com.ibm.uddi.datatypes` below for more information.)

### Enabling an EJB Client

This section is written on the assumption that WebSphere Application Server V5.0, a supported database and the IBM WebSphere UDDI Registry have already been installed.

#### Classpaths

In order for EJB clients to work correctly, the following jar files and folders must be added to the user's CLASSPATH:

#### For Windows

```
<WebSphere-install-dir>\lib\j2ee.jar  
<WebSphere-install-dir>\lib\naming.jar  
<WebSphere-install-dir>\lib\namingclient.jar  
<WebSphere-install-dir>\lib\ecutils.jar  
<WebSphere-install-dir>\lib\sas.jar  
<WebSphere-install-dir>\properties
```

#### For Unix Platforms, including z/OS

```
<WebSphere-install-dir>/lib/j2ee.jar  
<WebSphere-install-dir>/lib/naming.jar  
<WebSphere-install-dir>/lib/namingclient.jar  
<WebSphere-install-dir>/lib/ecutils.jar  
<WebSphere-install-dir>/lib/sas.jar  
<WebSphere-install-dir>/properties
```

In addition to these jars, there is also the jar file that contains all of the UDDI specific API for the EJB interface, which can be found at:

### For Windows

```
<DeploymentManager-install-dir>\UDDIReg\ejb\uddiejbclient.jar
```

where <DeploymentManager-install-dir> is the install location for WebSphere Application Server for Network Deployment, which by default is C:\Progra~1\WebSphere\DeploymentManager.

### For Unix Platforms, including z/OS

```
<DeploymentManager-install-dir>/UDDIReg/ebj/uddiejbclient.jar
```

where <DeploymentManager-install-dir> is the install location for WebSphere Application Server for Network Deployment, which by default is */opt/WebSphere/DeploymentManager* for Linux/Solaris systems or */usr/WebSphere/DeploymentManager* for AIX systems.

### The Path

Please ensure that your PATH statement starts with <WebSphere-install-dir>\java\bin

### Creating an EJB Client

If you want to read about creating EJB Clients in more detail, then please read the "Sun Microsystems Enterprise JavaBeans™ Specification Version 2.0"

### Finding the EJB Reference

An EJB Client can be a stand-alone Java application, an Applet, Servlet or a JSP. This document only covers writing a stand-alone Java application. In order to invoke an EJB that has been deployed into WebSphere on the server side, the Client must do two things: find the EJB on the server, and then create a Client side reference to that EJB. Once this Client side reference has been created, then the Client can invoke methods upon the EJB as if it was a local object. Clients cannot reference, or invoke, and EJB directly. Any calls made to the EJB must be made through the interfaces that the EJB provides. The interface that is used to create a local reference to the EJB is called the *home interface*. When an EJB is deployed in WebSphere, this home interface is made available to Clients by means of a searchable namespace. This means that a Client can look up an address on the namespace. If there is a home interface at that address, and it is the home interface to the EJB that they were looking for, then the Client can create a local instance of that home interface, and then, from that, a local reference to the EJB can be created.

### What code is needed in the Client?

The following code fragment illustrates how to Find and Create a local instance of the Inquiry EJB only. The same will need to be done to Find and Create a local copy of the Publish EJB.

```
private com.ibm.uddi.ejb.Inquiry inquiry = null;  
// This private variable, "inquiry" is going to be the local reference to the  
// EJB in WebSphere declaring it outside the scope of a method means that this
```

```

// same reference can be used throughout the client, without having to query
// the namespace again.
public void homeLookup()
{
    // These variables simply determine the address of the JNDI namespace, and
    // the address of the home interface within that namespace.

    // String naming_factory = "com.ibm.ejs.ns.jndi.CNInitialContextFactory";
    //WAS 4.0.2 Naming Factory
    String naming_factory = "com.ibm.websphere.naming.WsnInitialContextFactory";
    //WAS 5.0 Naming Factory

    String namespace_address = "iiop://localhost:2809/";
    //The address of the namespace
    String home_address = "com/ibm/uddi/ejb/InquiryHome";
    //The address of the home interface within the JNDI namespace

    java.util.Hashtable environment = new java.util.Hashtable();
    environment.put(javax.naming.Context.INITIAL_CONTEXT_FACTORY, naming_factory);
    environment.put(javax.naming.Context.PROVIDER_URL, namespace_address);

    try
    {
        javax.naming.InitialContext ic = new javax.naming.InitialContext(environment);
        // Create a context using the details above to connect to the namespace

        Object o = ic.lookup(home_address);
        // Do a lookup to see if there is an ejb_home at the address specified above

        // Now create a valid home instance for the EJB type we want to create
        com.ibm.uddi.ejb.InquiryHome home =
            (com.ibm.uddi.ejb.InquiryHome)(javax.rmi.PortableRemoteObject.narrow(o,
                com.ibm.uddi.ejb.InquiryHome.class));

        inquiry = home.create();
        // Now create a local reference of the EJB, by using the
        // home.create() method. Any business method that is intended for
        // the EJB in Websphere must be invoked against this inquiry object.
    }
    catch (javax.naming.NamingException ne) {ne.printStackTrace();}
    // This is thrown if there was a problem connecting to the namespace, or
    // finding the home_address in the namespace
    catch (java.rmi.RemoteException re) {re.printStackTrace();}
    // This usually indicates some sort of system failure, either WebSphere is
    // not running, or there is a communications problem
    catch (javax.ejb.CreateException ce) {ce.printStackTrace();}
    // This is thrown if the EJB reference cannot be created from the home instance.
}

```

### Writing Client code to use the EJB API

Once the reference to the EJB has been created (the Inquiry Object, in the above code), then the reference can be treated like any other Java Object. This is an example method using the UDDI EJB API - the only important point to remember is that, although the Inquiry Object has been created as a local reference, it is still referring to a remote EJB Object in a different server, possibly even in a different country. This means that at the very least a `javax.rmi.RemoteException` must be caught on each method call that is made to the EJB.

```

public void findBusiness()
{
    System.out.println("Find Business:");
    NameList names = new NameList();
    names.add(new Name("IBM Corporation"));
    //Create the list of names to find in the UDDI Registry, here just one
    // is used, "IBM Corporation"
}

```



```

    try
    {
        BusinessList list = inquiry.findBusiness(names);
//This is the call to the inquiry EJB that searches through the UDDI Registry

//Now display the amount of business found, and for each one, get the
// BusinessKey, the BusinessName and the amount of Services that Business has
        System.out.println("There are "+list.getBusinessInfos().size()+"
            matching Businesses in this registry");
        for (int i=0;i<list.getBusinessInfos().size();i++)
        {
            BusinessInfo business = list.getBusinessInfos().get(i);
            System.out.println("\nBusinessKey = "+business.getBusinessKey());
            System.out.println("BusinessName = "+business.getNames().get(0).getNameString());
            System.out.println("This Business Has "+business.getServiceInfos().size()+" Services\n");
        }
    }

// This is a UDDI specific exception, and will be thrown if for example an
// invalid name was used as the search criteria
catch (com.ibm.uddi.datatypes.DispositionReportException e) {this.handleDispositionReportException(e);}

        catch (java.rmi.RemoteException re) {re.printStackTrace();}
// This is the RemoteException that is thrown if there has been a system failure
// or a connection problem.
}

```

### What new code is needed on the Client?

Just as each EJB has an interface listed on the JNDI namespace, the `javax.transaction.UserTransaction` class also has an interface listed. This means that the same method used to get a local instance of an EJB can be applied to get a local instance of the `UserTransaction` class. Again, this code can be used to find the `UserTransaction` reference on the namespace, in addition to the code required to find the `Inquiry` EJB and the `Publish` EJB, or, alternatively, there is a slightly more elegant method used in the `TransactionEJBClientSample.java`.

```

public void txLookup()
{
    private javax.transaction.UserTransaction tx = null;
// This is the private variable that will be used to hold the UserTransaction Objec
// declaring it outside the scope of a method means that this same reference can be
// used throughout the client, without having to query the namespace again.

// These variables simply determine the address of the JNDI namespace, and the address
// of the home interface within that namespace.

// String naming_factory = "com.ibm.ejs.ns.jndi.CNInitialContextFactory";
//WAS 4.0.2 Naming Factory
String naming_factory = "com.ibm.websphere.naming.WsnInitialContextFactory";
//WAS 5.0 Naming Factory

String namespace_address = "iiop://localhost:2809/";
//The address of the namespace
String transaction_address = "jta/usertransaction";
//The address of the UserTransaction interface within the JNDI namespace

java.util.Hashtable environment = new java.util.Hashtable();
environment.put(javax.naming.Context.INITIAL_CONTEXT_FACTORY, naming_factory);
environment.put(javax.naming.Context.PROVIDER_URL, namespace_address);

try
{
    javax.naming.InitialContext ic = new javax.naming.InitialContext(environment);

```

```

// Create a context using the details above to connect to the namespace
    Object remote_object = ic.lookup(transaction_address);
// Do a lookup to see if there is a UserTransaction Object at the address specified above
tx = (javax.transaction.UserTransaction)remote_object;
//Convert the remote object found into a UserTransaction Object, and assign to the private variable
}
catch (javax.naming.NamingException ne) {ne.printStackTrace();}
// This is thrown if there was a problem connecting to the namespace, or
// finding the transaction_address in the namespace
}

```

### Writing Client code to use the EJB API with a Client transaction

To perform an Inquiry, a Publish or a Delete upon the IBM WebSphere UDDI Registry with client side transactional support requires very little additional code compared to doing the same operations without client side transactional support. Using the same code that is listed above (in "Writing Client Code to use the EJB API"), this example illustrates how easy client side transactions are to implement.

The additional lines of code needed are in bold type. This code also assumes that there is a variable called *tx* that has been declared at the class scope.

```

public void findBusiness()
{
//Just as there are UDDI and RMI specific exceptions thrown,
// 5 more exceptions need to be caught.
try
{
tx.begin(); //This begins the transaction context
System.out.println("Find Business:");
NameList names = new NameList();
names.add(new Name("IBM Corporation"));
//Create the list of names to find in the UDDI Registry, here just one is used, "IBM Corporation"
try
{
BusinessList list = inquiry.findBusiness(names);
//This is the call to the inquiry EJB that searches through the UDDI Registry

//Now display the amount of business found, and for each one, get the BusinessKey,
// the BusinessName and the amount of Services that Business has
System.out.println("There are "+list.getBusinessInfos().size()+" matching Businesses in this registry");
for (int i=0;i<list.getBusinessInfos().size();i++)
{
BusinessInfo business = list.getBusinessInfos().get(i);
System.out.println("\nBusinessKey = "+business.getBusinessKey());
System.out.println("BusinessName = "+business.getNames().get(0).getNameString());
System.out.println("This Business Has "+business.getServiceInfos().size()+" Services\n");
}
}
// This is a UDDI specific exception, and will be thrown if for example an invalid
// name was used as the search criteria
catch (com.ibm.uddi.datatypes.DispositionReportException e) {this.handleDispositionReportException(e);}
catch (java.rmi.RemoteException re) {re.printStackTrace();}
// This is the RemoteException that is thrown if there has been a system failure
// or a connection problem.

tx.commit(); //This ends the transaction context
}
catch (javax.transaction.NotSupportedException nse) {nse.printStackTrace();}
catch (javax.transaction.RollbackException rbe) {rbe.printStackTrace();}
catch (javax.transaction.SystemException se) {se.printStackTrace();}
catch (javax.transaction.HeuristicMixedException hme) {hme.printStackTrace();}
catch (javax.transaction.HeuristicRollbackException hrbe) {hrbe.printStackTrace();}
}

```

- "Datatypes package in the UDDI Registry"

- "EJB interface methods in the UDDI Registry"

## Datatypes package in the UDDI Registry

Below is a table listing the classes in the com.ibm.uddi.datatypes package, the elements in the UDDI v2.0 XML schema, and the correspondence between the two.

| com.ibm.uddi.datatypes Class |        | Corresponding UDDIv2.0 XML Schema Element      | Notes on DatatypeClass                                |
|------------------------------|--------|--|---|
| AccessPoint                  |        | accessPoint                                    |   |
| Address                      |        | address  |   |
| AddressLineList              | String | addressLine                                    | Encapsulates a Vector of addressLine Strings          |
| AddressList                  |        |  | Encapsulates a Vector of Address objects              |
| AssertionStatusItem          |        | assertionStatusItem                            |   |
| AssertionStatusItemList      |        |  | Encapsulates a Vector of AssertionStatusItem objects  |
| AssertionStatusReport        |        | assertionStatusReport (response message)       |   |
| AuthToken                    | String | authInfo                                       | Object containing authInfo String and operator name   |
| BindingDetail                | String | bindingKey<br>bindingDetail (response message) |   |
| BindingTemplate              |        | bindingTemplate                                |   |
| BindingTemplateList          |        | bindingTemplates                               | Encapsulates a Vector of Bindingtemplate objects      |
| BusinessDetail               |        | businessDetail (response message)              |   |
| BusinessDetailExt            |        | businessDetailExt (Response message)           | **  |
| BusinessEntity               |        | businessEntity                                 |   |
| BusinessEntityExt            |        | businessEntityExt                              | **  |
| BusinessEntityExtList        |        |  | Encapsulates a Vector of BusinessEntityExt objects ** |
| BusinessEntityList           |        |  | Encapsulates a Vector of BusinessEntity objects       |
| BusinessInfo                 |        | businessInfo                                   |   |
| BusinessInfoList             |        | businessInfo                                   | Encapsulates a Vector of businessInfo objects         |
| BusinessList                 | String | businessKey<br>businessList (response message) |   |
| BusinessService              |        | businessService                                |   |
| BusinessServiceList          |        | businessServices                               | Encapsulates a Vector of BusinessService objects      |
| CategoryBag                  |        | categoryBag                                    |   |
| Contact                      | String | completionStatus<br>contact                    |   |
| ContactList                  |        | contacts                                       | Encapsulates a Vector of Contact objects              |
| Description                  |        | description                                    |   |
| DescriptionList              |        |  | Encapsulates a Vector of Description objects          |
| DiscoveryUrl                 |        | discoveryURL                                   |   |

|                            |  |   |
|----------------------------|--|---|
| DiscoveryUrlList           | discoveryURLs                          | Encapsulates a Vector of DiscoveryURL objects   |
| DispositionReport          | dispositionReport                      | Exception thrown by EJB interface functions when an error occurs  |
| DispositionreportException |  |   |
| Email                      | email                                  | Encapsulates a Vector of Email objects  |
| EmailList                  |  |   |
| EndPoint                   |  | Used as baseclass for AccessPoint and HostingRedirector providing mutual exclusivity                                      |
| ErrInfo                    | errInfo                                | Encapsulates a FindQualifiers object and a maxrows field. Used in find_* API calls to specify search options              |
| FindQualifier              | findQualifier                          |   |
|                            | findQualifiers                         |   |
|                            | String fromKey                         |   |
| HostingRedirector          | hostingRedirector                      |   |
| IdentifierBag              | identifierbag                          |   |
| InquiryOptions             |  |   |
| InstanceDetails            | instanceDetails                        |   |
|                            | String instanceParms                   |   |
|                            | String keyValue                        |   |
| KeyedReference             | keyedReference                         | Abstract class, extended by some of the datatypes, which represents a string that can optionally be tagged with xml:lang. |
| keysOwned                  | keysOwned                              |   |
| LanguageString             |  |   |
| Name                       | name                                   | Encapsulates a Vector of Name objects   |
| NameList                   |  |   |
| OverviewDoc                | overviewDoc                            |   |
|                            | String overviewURL                     |   |
|                            | String personName                      |   |
| Phone                      | phone                                  | Encapsulates a Vector of Phone objects  |
| PhoneList                  |  |   |
| PublisherAssertion         | publisherAssertion                     | Encapsulates a Vector of Publisher Assertion objects  |
| PublisherAssertionList     |  |   |
| PublisherAssertions        | publisherAssertions (response message) |   |
| RegisteredInfo             | registeredInfo (response message)      |   |
| RelatesBusinessesList      | relatedBusinessInfo                    | not used  |
|                            | relatedBusinessInfos                   | not used  |
|                            | relatedBusinessesList                  |   |
| RelatedBusinessInfo        | relatedBusinessInfo                    |   |
| RelatedBusinessInfos       | relatedBusinessInfos                   |   |
| Result                     | result                                 |   |
| ResultList                 |  | Encapsulates a Vector of Result objects   |

|                        |        |                                  |   |
|------------------------|--------|----------------------------------|---|
| ServiceDetail          |        | serviceDetail (response message) |   |
| ServiceInfo            |        | serviceInfo                      |   |
| ServiceInfoList        |        | serviceInfos                     | Encapsulates a Vector of serviceInfo objects        |
|                        | String | serviceKey                       |   |
| ServiceList            |        | serviceList (response message)   |   |
|                        |        | sharedRelationships              | not used  |
| SharedRelationships    |        | sharedRelationships              |   |
| Tmodel                 |        | tModel                           |   |
| TModelBag              |        | tModelBag                        |   |
| TModelDetail           |        | tModelDetail (response message)  |   |
| TModelInfo             |        | tModelInfo                       |   |
| TModelInfoList         |        | tModelInfos                      | Encapsulates a Vector of TModelInfo objects         |
| TModelInstanceInfo     |        | tModelInstanceInfo               |   |
| TModelInstanceInfoList |        | tModelInstanceDetails            | Encapsulates a Vector of TModelInstanceInfo objects |
|                        | String | tModelKey                        |   |
| TModelList             |        | tModelList (response message)    |   |
| TModels                |        |                                  | Encapsulates a Vector of TModel objects             |
|                        | String | toKey                            |   |
|                        | String | uploadRegister                   |   |
| UploadRegisterList     |        |                                  | Encapsulates a Vector of uploadRegister strings     |

\*\* Used in UDDI API functions relating to BusinessDetailExtension. These UDDI API functions are not implemented in Version 1 of the IBM WebSphere UDDI Registry.

In general, a datatype called *DatatypeList* contains a vector of *Datatype* objects. Often these correspond to XML schema elements with plural names. (For example the datatype *Contact* corresponds to XML element *contact*, and *ContactList* corresponds to *contacts*.) Where there is no "plural" XML schema element for a particular *Datatype*, often there is still a *DatatypeList* where it is useful to have one, e.g. *AddressList*.

The exceptions to this naming convention occur when there is an existing XML schema element ending in "List". The exceptions are: *TModelList*, *ServiceList*, *BusinessList*. In these cases, the corresponding datatypes are given the same names as the XML schema elements, and the datatypes that would have had these names are called: *TModels*, *BusinessServiceList*, *BusinessEntityList*.

## EJB interface methods in the UDDI Registry

### Inquiry

```

findBinding
findBusiness
findRelatedBusinesses
findService
findTModel

```

```
getBindingDetail  
getBusinessDetail  
getServiceDetail  
getTModelDetail
```

### **Publish**

```
addPublisherAssertions  
deleteBinding  
deleteBusiness  
deletePublisherAssertions  
deleteService  
deleteTModel  
getAssertionStatusReport  
getRegisteredInfo  
getPublisherASsertions  
saveBinding  
saveBusiness  
saveService  
saveTModel  
setPublisherAssertions
```

Each method is overloaded and can take various combinations of arguments. The Javadoc contains detailed information about each method.

Note that `get_authToken` and `discard_authToken` are not implemented, as WebSphere security is used instead.

---

## **UDDI troubleshooting tips**

When the IBM WebSphere UDDI Registry is running, it might issue messages to report events or errors. You can use these messages, described in Messages as your first aid to problem determination. If you need more details about the causes of a problem, you can turn on tracing for UDDI, as described in Turning on UDDI trace.

- "Turning on UDDI trace"

Below are a few of the common causes of errors that might be found and their suggested solutions.

- If you set up the UDDI Registry application with a JDBC driver and datasource that reference Cloudscape, but set the `persist` property in `uddi.properties` to specify DB2, **or vice versa**, then some unexpected behavior will result, such as a fatal error on deleting an entity. If this happens, you should check that the above details are not in conflict. This only applies to a UDDI Registry installation on a single appserver.
- If you get a message "The application failed to initialize" when trying to access the UDDI User console and you are using DB2 as the persistence store for the UDDI Registry, a likely cause of the problem is that you specified the wrong `userid` and/or `password` when you ran the script to install the UDDI Registry application. If this occurs rerun the script ensuring you use the correct `userid` and `password`.
- You might find that, after uninstalling and reinstalling the UDDI Registry, you get errors from the UDDI User Console of the form:  
"Error 500: JSPG0059E: Unable to compile class for JSP".  
If this occurs, then you should clear out the temp directory of the WebSphere AppServer.
- When running one of the UDDI setup scripts `setupuddi.jacl` or `removeuddi.jacl`, if you get an error such as:

WASX7017E: Exception received while running file "setupuddi.jacl"; exception information: com.ibm.bsf.BSFException: error while eval'ing Jacl expression: java.util.MissingResourceException: Can't find resource for bundle java.util.PropertyResourceBundle, key ErrMsgIncorrectNumArgs

then please ensure that the file setupuddimessages.jafr is located in the *lib* subdirectory of the WebSphere deployment manager or application server under which you are running the script.

- When running the DB2 Setup Wizard, if you get an error stating "Invalid userid and password", then it could be caused by any of the following situations:
  - You have supplied an invalid userid or password - re-enter with a valid userid and password.
  - The supplied userid does not have the necessary privileges - retry with a userid that has appropriate privileges.
  - DB2 is stopped when you run the Wizard - start DB2 and retry the Wizard.
  - The UDDI20 database already exists and has been removed previously and, as such, is not catalogued. The DB2 Wizard does not recognize this situation and gives the error. You now have two options.
    1. If you wish to use the existing database then you will need to catalogue it and there is no need to rerun the Wizard.
    2. If you wish to create a new database you will need to recatalog the database and re-run the DB2 Wizard and choose the option to overwrite the database. (Any existing data **WILL** be lost.)

**Applicability of the following list: [Fix Pack 5.0.1 and later]**

- Catalog the database by:

- 

**Windows:**

```
>db2cmd
>db2cat -d iddi20
```

- 

**Unix platforms**

```
>su db2inst1 (or name of your db2 instance)
>db2 CATALOG DATABASE UDDI20
```

- **5.0.1** ensure that, if you are using a non-English installation of DB2, you have applied PTF1.

**5.0.1 Note:** If you have a copy of the file SetupDB2UDDI.jar in your appserver directory, then the application of the base and Network Deployment PTFs will not update SetupDB2UDDI.jar in your appserver directory. You must apply the PTF for Network Deployment to your DeploymentManager file structure to update the SetupDB2UDDI.jar located there (in the /IDDIReg/scripts subdirectory), and then manually copy this jar to any appserver you may wish to run it on.

- There is a limitation concerning URL rewriting causing JavaScript syntax errors on several Web pages in the UDDI User Console. Because of this, cookies must be enabled in client browsers, the application server must have cookies enabled as the session tracking mechanism, and URL rewriting must be disabled.
- If you have an existing DB2 version of the UDDI Registry database, and you use the UDDI DB2 setup wizard to replace this database with a new one, and if the database is in use at the time that you run the UDDI DB2 setup wizard, then the existing database is not overwritten.

- When running the UDDI DB2 setup wizard, as part of the installation step "Setting up the UDDI Registry to use DB2 within a deployment manager cell" or "Setting up the UDDI Registry to use DB2 in a single appserver", in addition to running the `was_install\bin\setupcmdline.bat` directory, you should also type either `set PATH=%WAS_PATH%` (for Windows platforms) or `export PATH=/opt/WebSphere/AppServer/java/bin:$PATH` (for Unix platforms) to ensure that you have access to Java.
- UDDI user console "Page cannot be displayed" errors with Internet Explorer.  
If you use Internet Explorer with the option "Show friendly HTTP error messages" enabled and you have WebSphere Application Server security enabled (user ID and password authentication enabled), you might experience intermittent errors on the browser, such as "Page cannot be displayed", when navigating the UDDI user console. This might be particularly noticeable when accessing the publish actions.  
To avoid such errors, disable the "Show friendly HTTP error messages" option on Internet Explorer. This option is found under **Tools > Internet options > Advanced Tab > Browsing Section**
- When using SOAP or UDDI4J, it is sometimes necessary to call `setServiceKey("")` before saving your changes, except with the EJB interface where this might result in an error.
- There are known problems with inquiries issued against the UDDI Registry if IBM Cloudscape is used as the persistence store for the registry data. Certain complex inquiries might produce unexpected results, or could fail. If your application needs to make inquiries of this nature, consider using DB2 as the persistence store. Note: DB2 must be used for production purposes. The IBM Cloudscape support is only provided for development and test use.
- You might see errors if you specify requests which specify more than 5 category values, more than 5 identifier values, or more than 5 technical model (tModel) values.
- If you stop and restart the UDDI Registry application from the administrative console, and then try to access the Registry through the user console, an "Error 500 - object is not an instance of a declaring class" displays on the user console, and the error message "SRVE0026E" displays in the system log. You cannot access the UDDI Registry until you restart the WebSphere Application Server. To avoid this problem when restarting the UDDI Registry, you should set **Prefer WEB-INF classes** on the panel navigated to by the following steps: **Applications > Manage Applications > UDDIRegistry > Web Modules defined for this Application > gui.war**.
- It is possible that a scripting error displays when you are running the `wsadmin appserversetupuddi.jacl` command.  
During installation, if you see the following error at the end of running the `appserversetupuddi.jacl` command, you can safely ignore the error. It is recommended that you start (or stop and restart) the application server and then continue.  
Here is an example of the error:  
UDIN2041I: Starting UDDI application. UDIN8019E: startApplication command for appname caught exception Exc. Values are: appname=UDDIRegistry, Exc=com.ibm.ws.scripting.Scripting Exception: com.ibm.websphere.management.exception.Connector Exception: ADMC0009E: Failed to make the SOAP RPC call: invoke

## Turning on UDDI trace

Before you begin



You can enable UDDI-specific trace in the same way as you enable other tracing in the WebSphere Application Server.

The following is a list of trace strings that may be used:

- com.ibm.uddi.api
- com.ibm.uddi.config
- com.ibm.uddi.datatypes
- com.ibm.uddi.dom
- com.ibm.uddi.ejb
- com.ibm.uddi.exception
- com.ibm.uddi.exceptions
- com.ibm.uddi.gui
- com.ibm.uddi.gui.inquire
- com.ibm.uddi.gui.publish
- com.ibm.uddi.persistence
- com.ibm.uddi.persistence.jdbc
- com.ibm.uddi.persistence.jdbc.cloudscape
- com.ibm.uddi.persistence.jdbc.db2
- com.ibm.uddi.ras
- com.ibm.uddi.security
- com.ibm.uddi.soap
- com.ibm.uddi.uuid
- com.ibm.uddi.validation
- com.ibm.uddi.xml

For example, to trace the UDDI User Console you would specify:

```
'com.ibm.uddi.gui=all=enabled'
```

This would enable all types of trace for the gui. Please refer to [ ] elsewhere in the WebSphere Application Server V5.0 InfoCenter for more information about using the administrator console to enable/disable trace.

---

## Messages

When the IBM WebSphere UDDI Registry is running, it might issue messages to report events or errors. The messages are in the form *UDxxnnnns* where:

**xx** is a two character descriptor identifying which component is involved

**nnnn** give the error code being issued

**s** is either I (Information) or E (Error)

The prefix *UDxxnnnns*: is followed by text which describes the event or error. For some messages, the first word of the text is one of the form (MSN=SSSS). The SSSS provides a message sequence number (or MSN), which identifies the unique circumstance in which the message was issued, and is of use where the same message can be issued in more than one circumstance.

To help you diagnose problems and minimize the need to enable trace in any of the above components, view the messages table. You can view the messages by

prefix or component, whichever is easiest for you to find in the table. All messages are documented with user/system action and explanation.

The text for the UDDI messages is contained in a file *uddiresourcebundles.jar* which is placed, by the installation process, into the `\lib` subdirectory (Windows) of the WebSphere application server into which the UDDI Registry was installed. If you will be running a console or log analyzer from another process; for example, to analyze the activity log, then you must place a copy of *uddiresourcebundles.jar* into a directory which is within the classpath of that process. Otherwise, the message lookup for the UDDI messages will fail.

### UDDI Components Message Prefix Table

click on individual links for message documentation for the component

|                               |   |
|-------------------------------|---|
| <a href="#rwsu_udai">UDAI</a> | <a href="#rwsu_udai">API</a>            |
| <a href="#rwsu_udcf">UDCF</a> | <a href="#rwsu_udcf">Configuration</a>  |
| <a href="#rwsu_udda">UDDA</a> | <a href="#rwsu_udda">Datatypes</a>      |
| <a href="#rwsu_uddm">UDDM</a> | <a href="#rwsu_uddm">DOM</a>            |
| <a href="#rwsu_udej">UDEJ</a> | <a href="#rwsu_udej">EJB Interface</a>  |
| <a href="#rwsu_udex">UDEX</a> | <a href="#rwsu_udex">Exceptions</a>     |
| <a href="#rwsu_udin">UDIN</a> | <a href="#rwsu_udin">Installation</a>   |
| <a href="#rwsu_udlc">UDLC</a> | <a href="#rwsu_udlc">Local API</a>      |
| <a href="#rwsu_udpr">UDPR</a> | <a href="#rwsu_udpr">Persistence</a>    |
| <a href="#rwsu_udrs">UDRS</a> | <a href="#rwsu_udrs">Logging</a>        |
| <a href="#rwsu_udsc">UDSC</a> | <a href="#rwsu_udsc">Security</a>       |
| <a href="#rwsu_udsp">UDSP</a> | <a href="#rwsu_udsp">SOAP Interface</a> |
| <a href="#rwsu_uduc">UDUC</a> | <a href="#rwsu_uduc">User Console</a>   |
| <a href="#rwsu_uduu">UDUU</a> | <a href="#rwsu_uduu">UUID</a>           |

## UDAI (Web Services UDDI) messages

There are no messages issued by this component.

## UDCF (Web Services UDDI) messages

**UDCF0001E: Exception occurred while getting int value of configuration property "<property>", exception: "<exception>"**

**Explanation:** This message is issued when an attempt to read the value of a configuration property from the `uddi.properties` file and convert it to integer has failed with the indicated exception.

**User Response:** Check that the `uddi.properties` file contains a value for the indicated configuration property, and that the value is valid. Check also that the indicated configuration property is a legal property. Refer to the InfoCenter for further information about global configuration properties and the `uddi.properties` file.

**UDCF0002E: Exception occurred while getting long value of configuration property "<property>", exception: "<exception>"**

**Explanation:** This message is issued when an attempt to read the value of a configuration property from the `uddi.properties` file and convert it to long has failed with the indicated exception.

**User Response:** Check that the `uddi.properties` file contains a value for the indicated configuration property, and that the value is valid. Check also that the indicated configuration property is a legal property. Refer to the InfoCenter for further information about global configuration properties and the `uddi.properties` file.

**UDCF0003E: Exception occurred while getting boolean value of configuration property "<property>", exception: "<exception>"**

**Explanation:** This message is issued when an attempt to read the value of a configuration property from the uddi.properties file and convert it to boolean has failed with the indicated exception

**User Response:** Check that the uddi.properties file contains a value for the indicated configuration property, and that the value is valid. Check also that the indicated configuration property is a legal property. Refer to the InfoCenter for further information about global configuration properties and the uddi.properties file.

**UDCF0004E: Failed to load UDDI global properties file.**

**Explanation:** This message is issued when the UDDI global configuration properties file, uddi.properties, cannot be loaded. Default values for the global configuration properties will be set, but these defaults may not be adequate for many of the properties, so you should investigate and resolve this problem.

**User Response:** Check that the uddi.properties file exists and is in the correct directory. Refer to the InfoCenter for further information about global configuration properties and the uddi.properties file.

**UDCF0005E: Exception occurred while loading UDDI global configuration properties, exception: "<exception>"**

**Explanation:** This message is issued when an attempt to load the UDDI global configuration properties from the uddi.properties has failed with the indicated exception. Default values for the global configuration properties will be set, but these defaults may not be adequate for many of the properties, so you should investigate and resolve this problem.

**User Response:** Check that the uddi.properties file exists and contains valid values for each of the configuration properties. Refer to the InfoCenter for further information about global configuration properties and the uddi.properties file.

## **UDDA (Web Services UDDI) messages**

There are no messages issued by this component.

## **UDDM (Web Services UDDI) messages**

There are no messages issued by this component.

## **UDEJ (Web Services UDDI) messages**

There are no messages issued by this component.

## **UDEX (Web Services UDDI) messages**

There are no messages issued by this component.

## **UDIN (Web Services UDDI) messages**

**UDIN0001I: Assuming hard coded defaults.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None..

**UDIN0002I: Cloudscape classpath is clpath. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0003I: Looking for childtype childname under parenttype parentname.**

**Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0004I: Looking for childtype childname under parenttype parentname and parenttype2 parentname2. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0005I: Conflict found with existing childtype childname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0006I: Not creating requested childtype. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0007I: Seeking parenttype with requested id of parentname. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0008I: Seeking parenttype with requested id of parentname under parenttype2 parentname2. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0009I: Attempting to create childtype under parenttype of parentID. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0010I: Create command that will be issued is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0011I: childtype childId was successfully created. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0012I: Looking for builtin\_rra.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0013I: List for J2CResourceAdapter returned N members. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0014I: Hunting J2CResourceAdapter associated with Node nodename.**

**Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0015I: Using rraID as builtin\_rra. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0016I: Using provider class of implclass with a classpath of clpath. Values**

**are:** **Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0017I: Installing to server servername, node nodename using database type of dbtype. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0018I: Attempting to create UDDI JDBCProvider.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0019I: Attempting to create UDDI Datasource.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0020I: Application Manager appmgr found. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0021I: Attempting to install UDDI Registry application.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0022I: Checking for installed UDDI Registry application of name appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

- UDIN0023W: Application of name appname is not present. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0024I: ApplicationManager not running, so application will not need to be stopped.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0025I: Stopping application of name appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0026W: stopApplication command for application appname caught exception Exc. Application might not have been running on this server. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0027I: Application appname stopped successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0028I: Removing application appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0029I: Application appname removed successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0030I: Adding resource bundles to repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0031I: Adding Cloudscape user functions to repository.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0032I: UDDI configuration properties file already exists. Only the persister and getServletURLPrefix properties will be overwritten.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0033I: Editing UDDI configuration properties file propsfile. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0034I: Url prefix found. Updating it to discoveryURL. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0035I: Persister property found. Updating it to dbtype. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0036I: Adding UDDI configuration properties file to repository for cell cellname under target node and server. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0037I: ws.ext.dir processing starting.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0038I: serverID is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0039I: JVM is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0040I: Out of N properties we located M matches at positions poslist.**

**Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0041I: Building new ws.ext.dirs properties.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0042I: SYSPROP is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0043I: ws.ext.dir has been set with new sysprop. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

- UDIN0044I: ws.ext.dir update skipped, required changes already present.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0045I: ws.ext.dir processing step complete.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0046I: Cleaning up temporary version of properties file temppropsfile.**  
**Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0047I: Issuing nodeSync.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0048I: UDDI Registry successfully installed. Please restart server servename to activate configuration changes. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0049I: Application Manager appmgr found. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0050I: Server is not running, so will not need to be stopped.**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0051I: Stopping server servename under node nodename. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0052I: Server servename stopped successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0053I: Restarting application server**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.  
**User Response:** None.
- UDIN0054I: Application server servename restarted successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.



**User Response:** None.

**UDIN0055I: Please ignore any errors concerning the serverStartupSyncEnabled attribute.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0101I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0102I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0103I: Changes were not saved on this call.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0104I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0105I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0106I: Attempting to save ws.ext.dir changes.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0107I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0108I: Attempting final save of new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN0109I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script setupuddi.jacl.

**User Response:** None.

**UDIN1001I: Application Manager appmgr found. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1002I: Server is not running, so will not need to be stopped.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1003I: Stopping server servername under node nodename. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1004I: Server servername stopped successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1005I: Resource bundles file will be removed from repository if present.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1006I: Removing resource bundles from repository.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1007I: Resource bundles successfully removed from repository.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1008I: Cloudscape user functions file will be removed from repository if present.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1009I: Removing Cloudscape user functions from repository.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1010I: Cloudscape user functions successfully removed from repository.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1011I: Application Manager appmgr found. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1012I: Checking for installed UDDI Registry application of name appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1013W: Application of name appname is not present. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1014I: ApplicationManager not running, so application will not need to be stopped.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1015I: Stopping application of name appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1016W: stopApplication command for application appname caught exception Exc. Application might not have been running on this server. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1017I: Application appname stopped successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1018I: Removing application appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1019I: Application appname removed successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1020I: UDDI datasource will be removed from server servername in node nodename if present. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1021I: Removing UDDI datasource.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1022I: UDDI datasource successfully removed.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1023I: UDDI JDBC driver will be removed from server servername in node nodename if present. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1024I: Removing UDDI JDBC driver.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1025I: UDDI JDBC driver successfully removed.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1026I: UDDI configuration properties file will be removed from repository if present.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1027I: Removing configuration properties file from cell cellname, node nodename and server servername. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1028I: Configuration properties file successfully removed from repository.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1029I: Issuing nodeSync.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1030I: UDDI Registry application, JDBC driver and datasource removed successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1031I: Restarting application server.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

- UDIN1032I: Application server servername restarted successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1033I: Please ignore any errors concerning the serverStartupSyncEnabled attribute.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1034I: ws.ext.dir processing starting.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1035I: serverID is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1036I: JVM is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1037I: Out of N properties we located M matches at positions poslist. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1038I: Removing UDDI values from ws.ext.dirs properties.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1039I: ws.ext.dir has been set with new sysprop. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1040I: ws.ext.dir update skipped, required changes already present.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1041I: ws.ext.dir processing step complete.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.  
**User Response:** None.
- UDIN1101I: Attempting to save new configuration.**  
**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1102I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1103I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1104I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1105I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1106I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1107I: Attempting final save of new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1108I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1109I: Attempting to save ws.ext.dir changes.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN1110I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script removeuddi.jacl.

**User Response:** None.

**UDIN2001I: Assuming hard coded defaults.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2002I: Listing members of type parenttype. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

- UDIN2003I: List for type parenttype returned N members. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2004I: Seeking parenttype with requested id of parentname. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2005I: Checking parentID with parentname. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2006I: Using this as parenttype of parentname. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2007I: Checking for existing childtype under parentname. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2008I: List for childtype returned N members. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2009I: No existing childtype present. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2010I: N existing objects of type childtype found, examining for conflict with childname. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2011I: Checking childID with name childname. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2012I: Conflict found with existing childtype of id childID. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2013I: Not creating requested object of type childtype. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.

- UDIN2014I: Conflict found with existing childtype, removing existing childtype. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2015I: Removal of childtype was successful. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2016I: Not in conflict.**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2017I: Attempting to create childtype under parentname of parentID. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2018I: Create command that will be issued is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2019I: childtype childID was successfully created. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2020I: No matches found.**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2021I: Looking for builtin\_rra.**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2022I: List for J2CResourceAdapter returned N members. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2023I: Hunting J2CResourceAdapter associated with Node nodename. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2024I: Using rraID as builtin\_rra. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.



**User Response:** None.

**UDIN2025I: Using provider class of implclass with a classpath of clpath. Values are:** **Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2026I: Installing to node nodename using database type of dbtype. Values are:** **Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2027I: Attempting to create UDDI JDBCProvider.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2028I: Attempting to create UDDI Datasource.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2029I: Application Manager appmgr found. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2030I: Attempting to install UDDI Registry application.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2031I: Checking for installed UDDI Registry application of name appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2032I: List for Applications returned N members. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2033W: Application of name appname is not present. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2034I: ApplicationManager not running, so application will not need to be stopped.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

- UDIN2035I: Stopping application of name appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2036W: stopApplication command for application appname caught exception Exc. Application might not have been running on this server. Values are:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2037I: Application appname stopped successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2038I: Removing application appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2039I: Application appname removed successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2040I: Attempting to install application appname. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2041I: Starting UDDI application.**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2042I: Application appname started successfully. Value is:**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2101I: Attempting to save new configuration.**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2102I: Changes saved successfully.**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.
- UDIN2103I: Changes were not saved on this call.**  
**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.  
**User Response:** None.

**UDIN2104I: Attempting to save post installation configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2105I: Changes saved successfully for UDDI Registry.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2106I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN2107I: Changes saved successfully for UDDI Registry.**

**Explanation:** This is an informational message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:** None.

**UDIN3001I: Application Manager appmgr found. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3002I: Checking for installed UDDI Registry application.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3003I: List for Applications returned N members. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3004W: Application of name appname is not present. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3005I: ApplicationManager not running, so application will not need to be stopped.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3006I: Stopping application of name appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3007W: stopApplication command for application appname caught exception Exc. Application might not have been running on this server. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3008I: Application appname stopped successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3009I: Removing application appname. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3010I: Application appname removed successfully. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3011I: UDDI datasource will be removed from server servername in node nodename if present. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3012I: Removing UDDI datasource.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3013I: UDDI datasource successfully removed.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3014I: UDDI JDBC driver will be removed from server servername in node nodename if present. Values are:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3015I: Removing UDDI JDBC driver from node nodename. Value is:**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3016I: UDDI JDBC driver successfully removed.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3017I: UDDI Registry application, JDBC driver and datasource removed successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

User Response: None.

**UDIN3101I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3102I: Changes to remove UDDI Registry have been saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3103I: Attempting to save new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3104I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3105I: Attempting final save of new configuration.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN3106I: Changes saved successfully.**

**Explanation:** This is an informational message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:** None.

**UDIN6001E: This script must be run in a Deployment Manager environment.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6002E: To install in a standalone application server, use appserversetupuddi.jacl.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6003E: Incorrect number of arguments passed to script.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6004E: Usage is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl. The text following 'Usage is:' gives the syntax for calling setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6005E: (<db2userid> <db2password> <db2ziplocation> are only required if setting up to use DB2).**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6006E: Use all forward (/) slashes to avoid problems with escaping back (\\) slashes.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6007E: Removal of childtype childname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6008E: An exception Exc occurred while creating childtype. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6009E: Unable to find requested parenttype of parentname. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6010E: List command for J2CResourceAdapter caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6011E: No J2CResourceAdapter objects available.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6012E: An error occurred during execution of setupuddi.jacl. Please check the parameters and try again.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6013E: Uninstall of application appname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6014E: Install of UDDI application caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6015E: Could not get JVM.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6016E: Cannot find nodeSync MBean.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6017E: nodeSync failed. UDDI Application may not be fully installed.**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6018E: stopServer command for server servername caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6019E: startServer command for server servername caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6101E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6102E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6103E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN6104E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script setupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7001E: This script must be run in a Deployment Manager environment.**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7002E: To remove from a standalone application server, use appserverremoveuddi.jacl.**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7003E: Incorrect number of arguments passed to script.**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7004E: Usage is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl. The text following 'Usage is:' gives the syntax for calling removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7005E: stopServer command for server servername caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7006E: Removal of resource bundles caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7007E: Removal of Cloudscape user functions caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7008E: Uninstall of application appname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7009E: Removal of UDDI datasource caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7010E: Removal of UDDI JDBC driver caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7011E: Removal of configuration properties file caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7012E: Cannot find nodeSync MBean.**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.



**UDIN7013E: nodeSync failed. UDDI Application may not be fully uninstalled.**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7014E: startServer command for server servername caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7015E: Could not get JVM.**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7101E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7102E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7103E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7104E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN7105E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script removeuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8001E: This script must be run on a standalone application server.**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8002E: To install in a Deployment Manager Environment, use setupuddi.jacl.**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8003E: Incorrect number of arguments passed to script.**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8004E: Usage is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl. The text following 'Usage is:' gives the syntax for calling appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8005E: (<db2userid> <db2password> <db2ziplocation> are only required if setting up to use DB2).**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8006E: Use all forward (/) slashes to avoid problems with escaping back (\\) slashes.**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8007E: List command for type parenttype caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8008E: No objects of type parenttype available. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8009E: List command for childtype caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8001E: This script must be run on a standalone application server.**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8010E: Error during remove of existing childtype, exception Exc caught. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8011E: Create command for childtype caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8012E: Unable to find requested parentype of parentname. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8013E: List command for J2CResourceAdapter caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8014E: No J2CResourceAdapter objects available.**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8015E: An error occurred during execution of appserversetupuddi.jacl. Please check the parameters and try again.**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8016E: List command for Applications caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8017E: Uninstall of application appname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8018E: Install of UDDI application caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8019E: startApplication command for appname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8101E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8102E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN8103E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserversetupuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9001E: This script must be run on a standalone application server.**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9002E: To remove from a deployment manager environment, use removeuddi.jacl.**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9003E: Incorrect number of arguments passed to script.**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9004E: Usage is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9005E: List command for Applications caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9006E: Uninstall of application appname caught exception Exc. Values are:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9007E: Removal of UDDI datasource caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9008E: Removal of UDDI JDBC driver caught exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9101E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9102E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

**UDIN9103E: Error saving configuration, changes not saved due to exception Exc. Value is:**

**Explanation:** This is an error message issued by the UDDI setup script appserverremoveuddi.jacl.

**User Response:**This message is self-explanatory.

## **UDLC (Web Services UDDI) messages**

There are no messages issued by this component.

## **UDPR (Web Services UDDI) messages**

There are no messages issued by this component.

## **UDRS (Web Services UDDI) messages**

**UDRS0001E: Exception "<exception>" occurred while attempting to get UDDI Message Logger.**

**Explanation:** This message is issued to stderr when an attempt to get the UDDI Message Logger fails with the indicated exception. Since the attempt to get the message logger failed, the message cannot be logged. No messages can be logged by this instance of the IBM WebSphere UDDI Registry.

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDRS0002E: Exception "<exception>" occurred while attempting to get UDDI Trace Logger for "<component>".**

**Explanation:** This message is logged when an attempt to get the UDDI Trace Logger for the specified component (or package) fails with the indicated exception. No trace entries can be logged for this component or package of the IBM WebSphere UDDI Registry.

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

## **UDSC (Web Services UDDI) messages**

There are no messages issued by this component.

## **UDSP (Web Services UDDI) messages**

**UDSP0001E: ParserPool found empty whilst attempting to process request. Request unsatisfied**

**Explanation:** A SOAP request was received, but was unable to be dealt with, as there were no free Parsers within the ParserPool.

**User Response:** Consider increasing the number of Parsers within the ParserPool by modifying the Init Parameter on the SOAP servlets.

**UDSP0002E: Error locating schemas required for UDDI processing. SOAP Servlets unworkable.**

**Explanation:** The SOAP servlet was unable to locate the schemas it requires in order to process SOAP requests. Without these, the servlet cannot process SOAP requests.

**User Response:** Check installation of UDDI was performed correctly. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0003W: Servlet unable to locate init parameter 'defaultPoolSize'. Using internal defaults.**

**Explanation:** The SOAP servlet was unable to locate the init parameter which sets the default size of the ParserPool. It will fall back to an internal default.

**User Response:** If this message occurred after attempting to make changes to the defaultPoolSize init parameter, ensure the changes were correct. If this message has appeared after installed, ensure installation was performed correctly.

**UDSP0004W: Servlet unable to understand init parameter 'defaultPoolSize'. Using internal defaults.**

**Explanation:** The SOAP servlet was unable to parse the init parameter which sets the default size of the ParserPool. It will fall back to an internal default.

**User Response:** If this message occurred after attempting to make changes to the defaultPoolSize init parameter, ensure the changes were correct. If this message has appeared after installed, ensure installation was performed correctly.

**UDSP0005E: Error occurred during parser creation.**

**Explanation:** An unspecified error occurred during the creation of a SOAP parser

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0006E: Internal configuration error.**

**Explanation:** This error may occur if there was a failure creating a Parser, with accompanying message UDSP0005. It may also occur if there was a problem acquiring the Persistence layer.

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0007E: Error during servlet acquisition of persistence layer.**

**Explanation:** The SOAP servlet was unable to acquire the persistence layer required for it to communicate with the UDDI datasource

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0008E: Error during servlet release of persistence layer.**

**Explanation:** The persistence layer reported a problem when the SOAP servlet attempted to release it.

**User Response:** Restart the UDDI registry. If the error persists, examine the WebSphere logs for information on its cause. If the problem cannot be resolved, then please contact the IBM Customer Service Center.

**UDSP0009E: Error during sending of response to client.**

**Explanation:** An error occurred when sending a SOAP response message back to a client. The client may not have received the response

**User Response:** This error is recorded to enable logging of failed responses to clients. The error may be the fault of the client disconnecting before the reply could be sent, or may indicate a network problem. Examine the WebSphere logs for more information on its cause.

## **UDUC (Web Services UDDI) messages**

**UDUC0001I: IBM WebSphere UDDI Registry user console starting initialization.**

**Explanation:** The user console control servlet is starting.

**User Response:** None.

**UDUC0002I: IBM WebSphere UDDI Registry user console finished initialization.**

**Explanation:** The user console control servlet has completed startup.

**User Response:** None.

**UDUC0003I: Reading init parameters.**

**Explanation:** The user console control servlet has started reading external parameters in its init method

**User Response:** None.

**UDUC0004I: Finished reading init parameters.**

**Explanation:** The user console control servlet has finished reading external parameters in its init method. This message indicates the user console is ready to accept client requests.

**User Response:** None.

**UDUC0005E: A serious error has occurred. Error message: <Message> error: <Throwable>. More information: <Additional information>.**

**Explanation:** This error message indicates an unexpected error has occurred. The <Message> describes the error that has occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** A trace of the gui component is recommended. Contact IBM support with this information.

**UDUC0006E: A persistence error has occurred. Error message: <Message> error: <Throwable>. More information: <Additional information>.**

**Explanation:** An error occurred while performing a database operation. The <Message> describes the error that occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** Check database connections and state. Please provide IBM support with a trace, including the gui and persistence components.

**UDUC0007E: A User mismatch error has occurred. Error message: <Message> error: <Throwable>. More information: <Additional information>.**

**Explanation:** The user id provided does not match the user id required or expected whilst performing an operation that requires authentication. The

<Message> describes the error that occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** Check the user has authority for the operation being requested. If necessary, contact IBM support detailing the actions taken to recreate the problem.

**UDUC0008E: An invalid key was passed. Error message: <Message> error: <Throwable>. More information: <Additional information>.**

**Explanation:** The requested operation is trying to retrieve information about an entity with a key that is invalid. This may occur if the entity has been deleted by another session. The <Message> describes the error that occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** Ask the client to close existing sessions and attempt the operation in a new browser session. If the problem persists, please provide IBM support with a trace of the gui and api components.

**UDUC0009E: An invalid value was supplied. Error message: <Message> error: <Throwable>. More information: <Additional information>**

**Explanation:** An invalid value was passed to an API call. The >Message> describes the error that occurred and the <Throwable> is the type of error that was caught. <Additional information> may provide further information, if available.

**User Response:** Contact IBM support with a trace of the gui and api components.

**UDUC0010E: Failed to introspect ActionForm properties. Exception: <Exception>.**

**Explanation:** String properties of a form object could not be introspected which means that the form contents cannot be checked for invalid characters.

**User Response:** Please contact IBM support with details of the Exception and a trace of the gui component.

**UDUC0011E: Failed to invoke reflected methods in ActionForm. Exception: <Exception>.**

**Explanation:** A form object's declared public method for setting or getting a String value could not be invoked. This method is required to check for invalid characters.

**User Response:** Please contact IBM support with details of the Exception and a trace of the gui component.

**UDUC0012E: User console initialization failed to connect to UDDI database. Exception: <Exception>.**

**Explanation:** During user console initialization, connection to the database failed, and threw the exception specified.

**User Response:** Check the connection to the UDDI database. The included exception message may yield some clues to help you resolve the problem. If unresolved, please contact IBM support with a trace of the gui component during startup.



**UDUC0013E: User console initialization failed to initialize tModels. Exception: <Exception>.**

**Explanation:** Indicates that an error has occurred during initialization of ActionServlet, specifically when reading tModels (invoking init method in class TModelNames).

**User Response:** Check the state of the UDDI database. Visually inspect the TMODEL table and confirm it is populated with valid data. The included exception message may yield some clues to help you resolve the problem. If unresolved, please contact IBM support with a trace of the gui component during startup.

**UDUC0014E: User console initialization failed to initialize taxonomies. Exception: <Exception>.**

**Explanation:** Indicates that an error has occurred during initialization of ActionServlet, specifically when reading taxonomy data (invoking init method of CategoryTaxonomyTree).

**User Response:** Check the state of the UDDI database. The included exception message may yield some clues to help you resolve the problem. If unresolved, please contact IBM support with a trace of the gui component during startup.

## UDUU (Web Services UDDI) messages

There are no messages issued by this component.

---

## Running the UDDI Samples

The UDDI samples, and documentation on how to use them, are available through the Web Services UDDI samples link on the Samples Central (<http://www.ibm.com/websphere/developer/library/samples/AppServer.html>) page of the IBM WebSphere Developer Domain Web site.

---

## Installation Verification Program (IVP)

### Before you begin

There are some samples available on the WSDD web site (at [ ]) that are intended to provide an optional Installation Verification test, or IVP, for the UDDI Registry component.

This topic describes how to run these installation verification programs (IVPs) to verify that the IBM UDDI Registry has been installed correctly.

There are two IVP SOAP samples called SOAPSsampleIVPa and SOAPSsampleIVPb. They are intended to verify the successful installation of the product, and should be used in conjunction with the UDDI Users Console (GUI). SOAPSsampleIVPa saves some data to the registry which you can then find using the GUI. Finally you can delete the data by running SOAPSsampleIVPb.

The IVP samples are installed into the same target directory as the other SOAP samples and they use the same XML files as the basic Java SOAP samples.

SOAPSsampleIVPa saves three businesses, six services (2 per business) and three tModels. The data structures are very basic and consist only of a name. The keys returned by the save\_\* UDDI API calls are then written to a file,

SOAPSamp1eIVPa.out. SOAPSamp1eIVPb then reads in these keys from the file in order to delete the saved data from the UDDI registry.

**Note:** Each time you run SOAPSamp1eIVPa, it overwrites the output file SOAPSamp1eIVPa.out so, if you wish to use SOAPSamp1eIVPb to delete the data, you must run this before you next run SOAPSamp1eIVPa.

To run the IVPs, complete the following steps on the same system as the UDDI Registry:

1. Ensure that DB2 and the WebSphere Admin Server are started.
2. Start the WebSphere Administrator's Console and ensure the default server is started and the UDDI Registry Application is started.
3. For SOAP samples to work, you need to ensure that the Client Developer Kit for Java is either the one shipped with IBM WebSphere Application Server or a later IBM Developer Kit for Java:
  - For Windows - ensure that <WebSphere-install-dir>\java\bin is present in the PATH statement before any other Developer Kits for Java
  - For Unix Platforms - ensure that <WebSphere-install-dir>/java/bin is present in the PATH statement before any other Developer Kits for Java

**Note:** You **must** use the IBM WebSphere supplied Developer Kit for Java or a later level of the IBM Developer Kit for Java.

For Windows, the default system path can be set via **Control Panel ...-> Settings ...-> System ...-> Advanced Properties ...-> Environment Variables**

Alternatively, this can be accomplished just for the shell where you plan to run the samples by modifying the path within the shell:

- For Windows - set path=<WebSphere-install-dir>\java\bin;%path%
  - For Unix Platforms - export PATH=<WebSphere-install-dir>/java/bin:\$PATH
4. Copy the samples and \*.xml files to a directory, and compile and run them there (see next bullets)
  5. Compile both SOAPSamp1eIVPa and SOAPSamp1eIVPb by typing:  
'javac SOAPSamp1eIVPa.java'  
  
and  
'javac SOAPSamp1eIVPb'  
  
.
  6. Run SOAPSamp1eIVPa by typing 'java SOAPSamp1eIVPa'. This should publish a number of businesses and services and technical models into the registry.
  7. Start your Web browser on the same system as the UDDI Registry.
  8. To display the UDDI GUI home page, type one of the following URLs:
    - If you have WebSphere security disabled: http://localhost:9080/uddigui
    - If you have WebSphere security enabled: https://localhost:9433/uddigui
  9. On the find page, complete the following steps:
    - a. Select the business radio button
    - b. In the data entry field, type % (percent is the wild card symbol)
    - c. Click **Find**

You should get a results page returned with three businesses (mybusiness1, mybusiness2, and mybusiness3). This demonstrates that the API and the UDDI Console are working correctly.

10. To see the services that are available for a business, click the "Show Services" option next to the business.
11. To delete all of the IVP data, run SOAPSampleIVPb (from the command prompt as before - by typing 'java SOAPSampleIVPb')
12. On the find page, complete the following steps:
  - a. Select the business radio button
  - b. In the data entry field, type % (percent is the wild card symbol)
  - c. Click **Find**

You should get an empty results page returned.

---

## Reporting problems with the IBM WebSphere UDDI Registry

### Before you begin

If you report a problem with the IBM WebSphere UDDI Registry component to IBM, please supply the following information:

1. A detailed description of the problem.
2. The build date and time of the version you are using. This can be obtained as follows:
  - In the installedApps subdirectory of the WebSphere install location, you will find a subdirectory called UDDI\_Registry.<nodename>.<servername>.ear, where <nodename> is the name of the node into which the UDDI Registry application is installed, and <servername> is the name of the server. Within that subdirectory, you will find a file called version.txt. Please include the contents of this file as part of your information.
  - If the UDDI Registry has been started with tracing enabled for the UDDI component, then you should find a trace entry in the WebSphere trace log which includes the strings " getUDDIMessageLogger" and "UDDI Build :"  
followed by the build date and time, and the build system. Please also include this information.
3. Any relevant log files and trace files.
  - If the problem occurred while setting up and installing the UDDI Registry application using one of the setup scripts, setupuddi.jacl or appserversetupuddi.jacl, then please supply the log output from running the script. (If you had not chosen to redirect the output from the script file to a log file, then please rerun the script, this time redirecting the output as described in the section 'Installing and Setting up a UDDI Registry'.) The log file will be in the directory from which you ran the setup script.
  - If the problem occurred while removing the UDDI Registry application using one of the remove scripts, removeuddi.jacl or appserverremoveuddi.jacl, then please supply the log output from running the script. (If you had not chosen to redirect the output from the script file to a log file, then please rerun the script, this time redirecting the output as described in the section 'Removing the UDDI Registry from a deployment manager cell' or 'Removing the UDDI Registry application from a single appserver'.) The log file will be in the directory from which you ran the remove script.
  - If the problem occurred while creating the UDDI Registry database using the UDDI DB2 Setup Wizard, then please supply the log file UDDIloadDB.log, which will be in the directory from which the wizard was run.

- If the problem occurred while running the UDDI Registry, please enable UDDI tracing (if not already enabled) and supply the trace log from the logs directory of the application server on which the UDDI Registry was running. Please refer to the section on 'Turning on UDDI Trace' for details on how to enable UDDI tracing.
4. If appropriate, any application code that you are using and the output produced by the application code.

In addition to the above, it might be useful to run the WebSphere collector tool and send the resulting jar file(s) (two files if run from Base AppServer AND DeploymentManager) to IBM. Instructions on running the collector tool can be found at [Running the collector tool](#)

---

## Feedback

### Before you begin

See the section on "Obtaining help from IBM" elsewhere in this InfoCenter for details on seeking assistance.

---

## Chapter 11. Web services gateway: Enabling Web services

You use the IBM Web services gateway to handle Web service invocations between Internet and Intranet environments. You use it to make your internal Web services available externally, and to make external Web services available to your internal systems.

With the Web services gateway you can administer Web services, Channels, Filters and UDDI registrations.

Detailed instructions on how to enable Web services through the IBM Web services gateway are given in the following tasks:

- Web services gateway - Completing the installation.
- Administering the Web services gateway.
- Running the Web services gateway samples.
- **5.0.2** Passing SOAP messages with attachments through the Web services gateway.
- Developing Web services gateway extensions.
- Administering security for the Web services gateway.
- Troubleshooting the Web services gateway.

**5.0.2** Before applying a WebSphere Application Server fix pack or upgrade, see Backing up and restoring a gateway configuration.

For a brief overview of what the Web services gateway is for, and how it works, see "Web services gateway - Frequently Asked Questions".

For a list of the major changes since the AlphaWorks preview version of the Web services gateway, see "Web services gateway - What is new in this release".

For additional technical details of the Web services gateway, see the gateway Javadoc(../javadoc/wsg/index.html).

For more information about working with Web services, visit the Internet sites referenced in Web services gateway: Resources for Learning.

---

### Web services gateway - Frequently Asked Questions

This topic provides answers to the following set of frequently asked questions about the Web services gateway:

- "What are Web services?" on page 648
- "What is the IBM Web services gateway?" on page 648
- "How does the Web services gateway work?" on page 648
- "What problems are solved by the Web services gateway?" on page 648
- "Who should use the Web services gateway?" on page 649
- **5.0.2** "What is the difference between the Apache SOAP channel and the SOAP/HTTP channel?" on page 649

## What are Web services?

Web services are modular applications that interact with one another across the Internet. Web services are based on shared, open and emerging technology standards and protocols (such as SOAP, UDDI, and WSDL) and can communicate, interact, and integrate with other applications, no matter how they are implemented.

## What is the IBM Web services gateway?

The gateway is a middleware component that bridges the gap between Internet and Intranet environments during Web service invocations. You use it to manage

- Web services.
- channels that carry requests to and responses from the services.
- filters that act upon the services.
- references to UDDI registries in which services can be registered.

## How does the Web services gateway work?

The gateway builds upon the Web Services Description Language (WSDL)(<http://www.w3.org/TR/wsdl>) and the ( Web Services Invocation Framework (WSIF)) for deployment and invocation.

You deploy a Web service to the Web services gateway by deploying a WSDL file that describes how the Web services gateway should access it. The WSDL file can be deployed to a UDDI registry or to a URL. You can send requests passing through the Web services gateway to a Java class, an enterprise bean, a SOAP server or a SOAP/JMS server (including another gateway).

A request to the Web services gateway arrives through a channel, is translated into an internal form, then passed through any filters that are registered for the requested service, and finally sent on to the service implementation. Responses follow the same path in reverse.

## What problems are solved by the Web services gateway?

- **Securely "externalizing" Web services:** Business applications that are exposed as Web services can be used by any Web service-enabled tool, regardless of the implementation details, to create new applications. To better integrate your business processes, you might want to expose these assets to business partners, customers and suppliers who are outside the firewall. The Web services gateway lets clients from outside the firewall use Web services that are buried deep inside your enterprise. The gateway also allows you to set access control on each of these deeply-buried services.
- **Better return on investment:** A process that you develop as a Web service can be reused by any number of partners.
- **Use of existing infrastructure:** With the Web services gateway, you can use your existing messaging infrastructure to make Web service requests, and use your existing Web services for external process integration.
- **Protocol transformation:** You might use one particular messaging protocol to invoke Web services, while your partners use some other protocol. Using the Web services gateway, you can trap the request from the client and transform it to another messaging protocol.

## Who should use the Web services gateway?

Any enterprise that chooses to share its resources selectively with its business partners and customers. IT managers and developers, who deploy resources, can also benefit from this technology.

## What is the difference between the Apache SOAP channel and the SOAP/HTTP channel?

### 5.0.2

The Apache SOAP channel and the SOAP/HTTP channel both support SOAP applications that are SOAP 1.1 compatible (for example Apache SOAP 2.3 and Axis SOAP 1.0). So if you have an application that uses a production-supported Axis 1.0 SOAP stack, generating SOAP 1.1, then this application can use either channel.

If you are using the Apache SOAP Channel, then the SOAP message format must be RPC style. To handle Document style SOAP messages, you must use the SOAP/HTTP channel (which supports both RPC style and Document style SOAP messages).

If you deploy Web services that pass attachments in a MIME message, then these Web services can only be accessed using the SOAP/HTTP channel.

---

## Web services gateway - What is new in this release

The Web services gateway was first made available on AlphaWorks(<http://www.alphaworks.ibm.com/tech/wsgw>) on 21 December 2001. The main differences between the AlphaWorks edition and this version are as follows:

- The gateway has been rebuilt using enterprise beans.  
**Note:** A side-effect of this is that the Web services gateway now only runs in an application server that has an EJB container. So it no longer runs in the Tomcat server.
- The gateway includes UDDI integration, so you can deploy and remove Web services to a UDDI registry as well as to a URL.
- **5.0.2** In addition to the Apache SOAP channel, there is now also a SOAP/HTTP channel.
- **5.0.2** The SOAP/HTTP channel supports
  - SOAP messages with attachments.
  - RPC style and Document style SOAP message formats.
- The gateway supports bidirectional interactions (that is, both inbound and outbound requests) directly, by deploying two instances of each type of channel.  
**Note:** To achieve this configuration with the AlphaWorks version, you had to deploy two instances of the Web services gateway; one for inbound communication and one for outbound communication.
- *Interceptors* have been renamed as *filters*.
- Channels, filters and UDDI references are deployed to the Web services gateway, then associated with individual Web services. So when you configure a Web service, you choose the following entities:
  - The channels on which it is available.

- The filters (if any) which apply to it.
- The UDDI references (if any) to which it is deployed.
- You can change the channels, filters and UDDI references that are associated with a deployed service without having to remove the service.
- You can deploy multiple targets for a single service (that is, more than one implementation of a service that has the same service interface).
- **5.0.2** You can back up and restore a gateway configuration.
- **5.0.2** You can create a gateway cluster (for example for load balancing).
- **5.0.2** You can configure the gateway for secure transmission of SOAP messages using tokens, keys, signatures and encryption in accordance with the emerging Web Services Security (WS-Security) specification(<http://www-106.ibm.com/developerworks/webservices/library/ws-secure/>).
- You can set security (basic authorization) on the individual methods of a Web service, and for the gateway as a whole.
- The gateway can invoke Web services over HTTPS.
- **5.0.2** You can configure the gateway to work with authenticating proxies.

---

## Web services gateway - Completing the installation

### Before you begin

This task assumes that, when you installed WebSphere Application Server, you either chose to install the Web services gateway (by choosing the "custom install" option **Web services -> Web services gateway**) or you accepted the "typical install" option (which includes the gateway). If you did this, then all the files that are needed to run a Web services gateway were copied into directories under *WebSphere\_DeployMgr\_root*, where *WebSphere\_DeployMgr\_root* is the deployment manager root directory (by default WebSphere/DeploymentManager).

The following table lists the Web services gateway files, and the locations into which they are placed by the install. The **Location** column shows the subdirectory under *WebSphere\_DeployMgr\_root*. For example, if you installed WebSphere Application Server onto a machine running Windows, and accepted the default directory names, then the location of the "installableApps" directory is C:\Progra~1\WebSphere\DeploymentManager\installableApps.

| File name         | Purpose   | Location         |
|-------------------|---|------------------|
| wsgw.ear          | The Web services gateway application                              | /installableApps |
| wsgwsoap1.ear     | The Apache SOAP channel application number 1                      | /installableApps |
| wsgwsoap2.ear     | The Apache SOAP channel application number 2                      | /installableApps |
| wsgwsoaphttp1.ear | The SOAP/HTTP channel application number 1                        | /installableApps |
| wsgwsoaphttp2.ear | The SOAP/HTTP channel application number 2                        | /installableApps |
| wsgwcorr.ear      | The application used to correlate any asynchronous reply messages | /installableApps |



|                           |   |                       |
|---------------------------|---|-----------------------|
| wsgwauth.ear              | The Web service operation-level security application  | /installableApps      |
| WSGWResourceBundles.jar   | System messages for the Web services gateway application  | /lib                  |
| Various install scripts   | Installation of the Web services gateway  | /WSGW/scripts/install |
| Authorization scripts     | Generation of authorization beans for Web service operation-level security  | /WSGW/scripts/auth    |
| Cloudscape directory tree | Cloudscape database containing the Web services gateway tables and pre-loaded data needed to install a database to use with the gateway | /bin/WSGWDB           |

**5.0.2** If you want to run a gateway cluster (for example to guarantee high availability) you must create the cluster before you install the gateway on any machine in the cluster.

To complete the gateway installation, you have two choices:

- Install into a deployment manager cell.
- Install into a stand-alone application server.

For production use you will probably want to install the gateway into one or more nodes of a deployment manager cell. But for development or test purposes you might find it useful to install the gateway into a stand-alone application server (that is, an application server that is not part of a deployment manager cell).

#### What to do next

To finish the Web services gateway installation, complete the following steps:

1. Confirm that your system configuration complies with the Web services gateway prerequisites and constraints.
2. Establish requirements for using a database with the gateway.
3. Enable security if required.
4. Either Install the gateway into a deployment manager cell, or Install the gateway into a stand-alone application server.
5. Test the installation.

## Web services gateway - prerequisites and constraints

The gateway is an enterprise application that you can deploy

- to an application server that is a node of a WebSphere Application Server Network Deployment cell.
- to a stand-alone application server.

**5.0.2** If you want to run a gateway cluster (for example to guarantee high availability) you must create the cluster (probably as a deployment manager cell) before you install the gateway on any machine in the cluster.

WebSphere Application Server Network Deployment is not a stand-alone product for running enterprise applications. So if you want to deploy the gateway you also need to install the base WebSphere Application Server product.

**Note:** Although it is not installed by default, a copy of the base WebSphere Application Server product is packaged with the WebSphere Application Server Network Deployment product.

This version of the Web services gateway is also subject to the following constraints:

- WSDL definitions for target services must use XML Schema version 2001. For more information, see "Web services gateway troubleshooting tips".

- **5.0+**

- **5.0.1+**

- The SOAP message format must be RPC style. The gateway does not support Document style SOAP messages.

- **5.0.2** If you are using the Apache SOAP Channel, then the SOAP message format must be RPC style. To handle Document style SOAP messages, use the SOAP/HTTP channel (which supports both RPC style and Document style SOAP messages).
- The gateway application (wsgw.ear) must be installed before channel and filter applications. If the gateway application needs to be reinstalled, all channels and filters must be uninstalled first, then reinstalled after the gateway application.
- **5.0+** The gateway does not support WSDL service definitions that contain **soap:header** elements within their **wSDL:definition** element.
- **5.0.1** If the WSDL for your service contains **soap:header** elements within the **wSDL:definition** element, then the gateway passes the SOAP headers through. This is correct behavior. However, you will also see the following effects:
  - The SOAP headers will not be included in the WSDL that the gateway generates.
  - If you set the "must understand" flag on the SOAP message, then you will get an error message.

You might also find it useful to enable trace for all gateway components: Set the WebSphere Application Server trace string to `com.ibm.wsgw.*=all=enabled`, and have trace, stdout and stderr for the application server written to a well-known location. For information on how to do this, see Enabling trace.

## Establishing requirements for using a database with the gateway

Before installing the Web services gateway you need to decide on your database requirements. There are three choices:

### No database

You do not have to install a database, but if you do not install one you cannot use asynchronous channels.

- DB2** If DB2 is already installed on your system, you can create an associated DB2 Web services gateway database.

## Cloudscape

You can use the copy of Cloudscape runtime that is installed with WebSphere Application Server.

### What to do next

Database installation instructions are included in the next task. So after you have decided on your database requirements you are ready to *either* Install the gateway into a deployment manager cell *or* Install the gateway into a stand-alone application server.

## Installing the gateway into a deployment manager cell

### Before you begin

This task assumes that you have already completed the step Establishing requirements for using a database with the gateway.

If you want to enable gateway-level security, you must do so before you install the Web services gateway.

If you want to use DB2 on a UNIX system, the shell in which the install is invoked requires the environment to be set up using the db2profile and usejdbc2 script files. To do this, run both the following commands:

```
./home/db2_instance/sql11ib/db2profile  
./home/db2_instance/sql11ib/java12/usejdbc2
```

where *db2\_instance* is the db2 instance to be configured.

In this task you install the gateway into an application server that is hosted on a node of an existing deployment manager cell. The major elements of this process are as follows:

- Database and Table creation (optional - DB2 only).
- JDBC driver and Datasource creation (optional - only if a database is being installed).
- Installation of the following enterprise applications:
  - The gateway application.
  - Apache SOAP channel 1.
  - **5.0.2** SOAP/HTTP channel 1.
  - The correlation application (optional - only if a database is being installed).
- Installation of other gateway applications. For example:
  - Apache SOAP channel 2.
  - **5.0.2** SOAP/HTTP channel 2.

To install the gateway into an application server that is hosted on a node of a deployment manager cell, complete the following steps:

### Steps for this task

1. **(Optional)** To create and install a DB2 database and associated tables, complete the following steps:
  - a. From a command prompt, go to directory `WebSphere_DeployMgr_root/WSGW/scripts/install`, where

*WebSphere\_DeployMgr\_root* is the deployment manager root directory (by default *WebSphere/DeploymentManager*).

b. Enter the command `WSGWinstallDB`:

- **[Windows]** `WSGWinstallDB.bat WebSphere_DeployMgr_root db2user_id db2password`
- **[UNIX]** `WSGWinstallDB.sh WebSphere_DeployMgr_root db2user_id`

For example (UNIX systems):

```
./WSGWinstallDB.sh /opt/WebSphere/DeploymentManager mydb2id
```

**Note:** Running `WSGWinstallDB` also creates `WSGWinstallDB.log` in the application server for network deployment's `/logs` directory. Open this file to check that the database was created successfully.

2. Start the application server (you can use the command `startServer.ext your_server`).

3. From a command prompt, go to directory `WebSphere_DeployMgr_root/WSGW/scripts/install`.

4. Clear your class path.

You can use the following command:

- (Windows systems): `set CLASSPATH=`
- (UNIX systems): `unset CLASSPATH`

5. Enter the command `WebSphere_DeployMgr_root/bin/wsadmin.ext -f setupWSGW.jacl parm1 ... parmN`

where

- *parm1* is the *WebSphere\_DeployMgr\_root* directory
- *parm2* is the server name
- *parm3* is the node name (this is case sensitive)

**If you are not using a database**, there are no more parameters and `wsgwcorr.ear` is not installed. For example (Windows systems):

```
C:\Progra~1\WebSphere\DeploymentManager\bin\wsadmin.bat -f  
setupWSGW.jacl C:/Progra~1/WebSphere/DeploymentManager server1 PHJ2
```

**Note:** The use of forward slashes (/) is compulsory for this command, even on Windows systems.

**If you are using Cloudscape**, there is one more parameter:

- *parm4* is the name and location of the `WSGWDB` directory (*WebSphere\_DeployMgr\_root/bin/WSGWDB*).

For example (Windows systems):

```
C:\Progra~1\WebSphere\DeploymentManager\bin\wsadmin.bat -f  
setupWSGW.jacl C:/Progra~1/WebSphere/DeploymentManager server1 PHJ2  
C:/Progra~1/WebSphere/DeploymentManager/bin/WSGWDB
```

**Note:** The use of forward slashes (/) is compulsory for this command, even on Windows systems.

**If you are using DB2**, there are four more parameters:

- *parm4* is WSGWDB.
- *parm5* is your DB2 user ID
- *parm6* is your DB2 password
- *parm7* is the name and location of file *db2java.zip*

For example (Windows systems):

```
C:\Progra~1\WebSphere\DeploymentManager\bin\wsadmin.bat -f
setupWSGW.jac1 C:/Progra~1/WebSphere/DeploymentManager server1 PHJ2
WSGWDB mydb2id mydb2pw C:/Progra~1/SQLLIB/java/db2java.zip
```

**Note:** The use of forward slashes (/) is compulsory for this command, even on Windows systems.

When you run the `setupWSGW.jac1` command, the following initial set of gateway applications is installed:

- The gateway itself (*wsgw.ear*).
  - Apache SOAP channel 1 (*wsgwsoap1.ear*).
  - **5.0.2** SOAP/HTTP channel 1 (*wsgwsoaphttp1.ear*).
  - (Optionally) The correlation application (*wsgwcorr.ear*).
6. Install additional gateway applications. For example:
- Apache SOAP channel 2 (*wsgwsoap2.ear*)
  - **5.0.2** SOAP/HTTP channel 2 (*wsgwsoaphttp2.ear*)

To install additional gateway applications, complete the following steps:

**Note:** If you prefer, you can install these EAR files using the WebSphere Application Server administrative console, as described in the final step of the task Enabling operation-level authorization.

- From a command prompt, go to directory *WebSphere\_DeployMgr\_root/bin*
- To start the application server, enter the command `startServer.ext` *your\_server*
- To start the WebSphere administration program, enter the command `wsadmin.ext`.
- For each additional Web services gateway enterprise application that you want to install, enter the following commands at the `wsadmin>` prompt:

```
$AdminApp install path_to_ear_file {-appname application
                                     -server your_server -node your_node_name}
$AdminConfig save
```

where

- *application* is the name of the enterprise application
- *path\_to\_ear\_file* is the name and location of the enterprise application's EAR file
- *your\_node\_name* is the node name (this is case sensitive)

For example (Windows systems):

```
wsadmin>$AdminApp install
      C:/Progra~1/WebSphere/DeploymentManager/installableApps/wsgwsoap2.ear
      {-appname wsgwsoap2 -server server1 -node PHJ2}
$AdminConfig save
```

- e. After you have installed all your additional Web services gateway enterprise applications, close the WebSphere administration program by entering `quit` or `exit` at the `wsadmin>` prompt.
7. To stop then restart the application server, complete the following steps:
  - a. Enter the command `stopServer.ext your_server`
  - b. Enter the command `startServer.ext your_server`

### What to do next

If you want more than one gateway installation (for example to create a gateway cluster) then repeat the steps given in this topic for another application server that is hosted on a node in a deployment manager cell.

You are now ready to test the installation. Run the test on every application server on which you have installed the gateway.

## Installing the gateway into a stand-alone application server

### Before you begin

This task assumes that you have already completed the step Establishing requirements for using a database with the gateway.

If you want to enable gateway-level security, you must do so before you install the gateway.

If you want to use DB2 on a UNIX system, the shell in which the `install` is invoked requires the environment to be set up using the `db2profile` and `usejdbc2` script files. To do this, run both the following commands:

```
./home/db2_instance/sqllib/db2profile  
./home/db2_instance/sqllib/java12/usejdbc2
```

where `db2_instance` is the db2 instance to be configured.

**Note:** The application server in which you run the gateway must not form part of a cell managed by a deployment manager. In other words, you must not issue an `addNode` command for the node containing the application server in which you run the Web services gateway application. If you do issue the `addNode` command, then the installed Web services gateway application is deleted by the node synchronization process that takes place within a cell of application servers.

When you installed WebSphere Application Server, all the files that are needed to run a Web services gateway were copied into directories under `WebSphere_DeployMgr_root`, where `WebSphere_DeployMgr_root` is the deployment manager root directory (by default `WebSphere/DeploymentManager`). Before you can install and run the gateway in a single application server instance in your network space, you must first copy these files over to the application server. You can do this by completing the following steps:

- Stop the application server into which you plan to install the Web services gateway. You can use the command `stopServer.ext your_server` where
  - `.ext` is `.bat` for a Windows system and `.sh` for a UNIX system.
  - `your_server` is your application server's name.

For example (UNIX systems): `./stopServer.sh server1`.

- Copy all the EAR files with filenames that begin "wsgw" from the *WebSphere\_DeployMgr\_root/installableApps* directory of the machine on which you installed WebSphere Application Server Network Deployment into the *stand-alone\_WAS\_HOME/installableApps* directory of the target application server's install tree, where *stand-alone\_WAS\_HOME* is the root directory of your target application server (by default WebSphere/AppServer).
- Copy file *WebSphere\_DeployMgr\_root/lib/WSGWResouceBundles.jar* into directory *stand-alone\_WAS\_HOME/lib*.
- Copy directory *WebSphere\_DeployMgr\_root/WSGW* and all files and directories within it into directory *stand-alone\_WAS\_HOME/WSGW*.
- If you plan to use the Cloudscape database with the Web services gateway, then copy directory *WebSphere\_DeployMgr\_root/bin/WSGWDB* and all files and directories within it into directory *stand-alone\_WAS\_HOME/bin/WSGWDB*.

In this task you install the gateway into an individual application server instance in your network space. The major elements of this process are as follows:

- Database and Table creation (optional - DB2 only).
- JDBC driver and Datasource creation (optional - only if a database is being installed).
- Installation of the following enterprise applications:
  - The gateway application.
  - Apache SOAP channel 1.
  - **5.0.2** SOAP/HTTP channel 1.
  - The correlation application (optional - only if a database is being installed).
- Installation of other gateway applications. For example:
  - Apache SOAP channel 2.
  - **5.0.2** SOAP/HTTP channel 2.

To install the gateway into an application server instance, complete the following steps:

#### Steps for this task

1. **(Optional)** To create and install a DB2 database and associated tables, complete the following steps:
  - a. From a command prompt, go to directory *stand-alone\_WAS\_HOME/WSGW/scripts/install*.
  - b. Enter the command `WSGWinstallDB`:
    - **[Windows]** `WSGWinstallDB.bat stand-alone_WAS_HOME db2user_id db2password`
    - **[UNIX]** `WSGWinstallDB.sh stand-alone_WAS_HOME db2user_id`

For example (UNIX systems):

```
./WSGWinstallDB.sh /opt/WebSphere/AppServer mydb2id
```

**Note:** Running `WSGWinstallDB` also creates `WSGWInstallDB.log` in the stand-alone application server's `/logs` directory. Open this file to check that the database was created successfully.

2. Start the application server (you can use the command `startServer.ext your_server`).

3. From a command prompt, go to directory *stand-alone\_WAS\_HOME*/WSGW/scripts/install.
4. Clear your class path.  
You can use the following command:
  - (Windows systems): set CLASSPATH=
  - (UNIX systems): unset CLASSPATH
5. Enter the command *stand-alone\_WAS\_HOME*/bin/wsadmin.ext -f setupWSGW.jacl *parm1* ... *parmN*  
where
  - *parm1* is the *stand-alone\_WAS\_HOME* directory
  - *parm2* is the server name
  - *parm3* is the node name (this is case sensitive)

**If you are not using a database**, there are no more parameters and wsgwcorr.ear is not installed. For example (Windows systems):

```
C:\Progra~1\WebSphere\AppServer\bin\wsadmin.bat -f setupWSGW.jacl
C:/Progra~1/WebSphere/AppServer server1 PHJ2
```

**Note:** The use of forward slashes (/) is compulsory for this command, even on Windows systems.

**If you are using Cloudscape**, there is one more parameter:

- *parm4* is the name and location of the WSGWDB directory (*stand-alone\_WAS\_HOME*/bin/WSGWDB).

For example (Windows systems):

```
C:\Progra~1\WebSphere\AppServer\bin\wsadmin.bat -f setupWSGW.jacl
C:/Progra~1/WebSphere/AppServer server1 PHJ2
C:/Progra~1/WebSphere/AppServer/bin/WSGWDB
```

**Note:** The use of forward slashes (/) is compulsory for this command, even on Windows systems.

**If you are using DB2**, there are four more parameters:

- *parm4* is WSGWDB.
- *parm5* is your DB2 user ID
- *parm6* is your DB2 password
- *parm7* is the name and location of file db2java.zip

For example (Windows systems):

```
C:\Progra~1\WebSphere\AppServer\bin\wsadmin.bat -f setupWSGW.jacl
C:/Progra~1/WebSphere/AppServer server1 PHJ2 WSGWDB mydb2id mydb2pw
C:/Progra~1/SQLLIB/java/db2java.zip
```

**Note:** The use of forward slashes (/) is compulsory for this command, even on Windows systems.

When you run the setupWSGW.jacl command, the following initial set of gateway applications is installed:



- The gateway itself (wsgw.ear).
  - Apache SOAP channel 1 (wsgwsoap1.ear).
  - **5.0.2** SOAP/HTTP channel 1 (wsgwsoaphttp1.ear).
  - (Optionally) The correlation application (wsgwcorr.ear).
6. Install additional gateway applications. For example:
- Apache SOAP channel 2 (wsgwsoap2.ear)
  - **5.0.2** SOAP/HTTP channel 2 (wsgwsoaphttp2.ear)

To install additional gateway applications, complete the following steps:

**Note:** If you prefer, you can install these EAR files using the WebSphere Application Server administrative console, as described in the final step of the task Enabling operation-level authorization.

- a. From a command prompt, go to directory *stand-alone\_WAS\_HOME/bin*
- b. To start the application server, enter the command `startServer.ext your_server`
- c. To start the WebSphere administration program, enter the command `wsadmin.ext`.
- d. For each additional Web services gateway enterprise application that you want to install, enter the following commands at the `wsadmin>` prompt:

```
$AdminApp install path_to_ear_file {-appname application
                                -server your_server -node your_node_name}
$AdminConfig save
```

where

- *application* is the name of the enterprise application
- *path\_to\_ear\_file* is the name and location of the enterprise application's EAR file
- *your\_node\_name* is the node name (this is case sensitive)

For example (Windows systems):

```
wsadmin>$AdminApp install
        C:/Progra~1/WebSphere/AppServer/installableApps/wsgwsoap2.ear
        {-appname wsgwsoap2 -server server1 -node PHJ2}
$AdminConfig save
```

- e. After you have installed all your additional Web services gateway enterprise applications, close the WebSphere administration program by entering `quit` or `exit` at the `wsadmin>` prompt.
7. To stop then restart the application server, complete the following steps:
- a. Enter the command `stopServer.ext your_server`
  - b. Enter the command `startServer.ext your_server`

### What to do next

You are now ready to test the installation.

## Testing the Web services gateway installation

Use this task to test that the Web services gateway has been installed correctly.

To test the basic installation of the Web services gateway, complete the following steps:

#### Steps for this task

1. In a Web browser, go to `http://host:port/wsgw` where *host* and *port* are the host name and port number that your HTTP server is listening on.  
The browser should display the following message:  
**IBM Web services gateway**  
What do you want to do?
  - Run the admin client
  - View the product ID
2. If the previous step was successful, then to test the Apache SOAP channels use your Web browser to display the Web page at `http://host:port/wsgwengine/soaprpcrouter` where *engine* is either *soap1* or *soap2*.  
The browser should display the following message: Sorry, I don't speak via HTTP GET - you have to use HTTP POST to talk to me.
3. If the previous step was successful, then complete the following steps:
  - a. To test the Apache SOAP channels, use your Web browser to display the Web page at `http://host:port/wsgwengine/soaprpcrouter` where *engine* is either *soap1* or *soap2*.  
The browser should display the following message: Sorry, I don't speak via HTTP GET - you have to use HTTP POST to talk to me.
  - b. To test the SOAP/HTTP channels, use your Web browser to display the Web page at `http://host:port/wsgwengine/soaphttpengine` where *engine* is either *soaphttp1* or *soaphttp2*.  
The browser should display the message: And now... Some Services, followed by a list of the Web services that are currently deployed to the channel.

If you do not see these messages, your server is not configured correctly - in which case, see Web services gateway troubleshooting tips.

---

## Backing up and restoring a gateway configuration

The **Back up** options write out, to a single file, the deployment details for all the channels, filters, UDDI references and Web services that are currently deployed to the gateway. The **Restore** option uses the information contained in a previously-saved file to populate an empty installation of the gateway with the same deployment details for those channels, filters, UDDI references and Web services.

There are two situations in which you might want to back up a gateway configuration:

- Before you apply an upgrade or fix pack to WebSphere Application Server.
- Because you want to create or update a gateway cluster.

When you *apply an upgrade or fix pack*, the configured gateway is replaced with an upgraded but empty gateway. You use the **Save Gateway Configuration -> Private** option to preserve your configuration.

When you *want to create or update a gateway cluster*, you use the **Save Gateway Configuration -> Shared** option to save a "shareable" version of your gateway configuration. This saved configuration does not include machine-specific values, and therefore does not over-write these values on the target gateways in your cluster when you restore the configuration to them.

The **restore** option automatically detects whether the configuration file that it is restoring contains a **Private** or a **Shared** gateway configuration.

**Note:** For the **Save Gateway Configuration -> Private** option, the steps below include

- Manual un-publishing of Web services from UDDI before backing up the gateway configuration.
- Manual re-publishing of these Web services to UDDI after restoring the gateway configuration.

The reasons for this are explained in Backing up and restoring UDDI publication links.

To back up and restore a gateway configuration, complete the following steps:

#### Steps for this task

1. Display the Web services gateway administrative user interface.
2. **(Optional)** If you are backing up before you apply an upgrade or fix pack, and any of your deployed Web services are also published by the gateway to UDDI registries, use the Listing and managing gateway-deployed Web services option to un-publish them from UDDI. Make a note of the UDDI deployment details so that you can re-publish them to UDDI after you restore the gateway configuration.
3. To **back up** the gateway configuration, complete the following steps:

a.

- 1) In the gateway user interface navigation pane, click the following link:

#### **Gateway**

- Back Up

The **Save Gateway Configuration** screen is displayed.

- 2) In the **Location** field, type the path to your configuration file.

#### **Note:**

- The path to your configuration file must point to a local drive on the machine on which the gateway is currently running.
  - The file name for the configuration file can be any valid Java filename.
- 3) If you are backing up before you apply an upgrade or fix pack, select **Private**. If you are backing up to create a gateway cluster, select **Shared**.
  - 4) Click **OK**

Your gateway configuration is saved to the location you specified.

- b. To **restore** a gateway configuration, complete the following steps:

- 1) Check that the target gateway is empty. If there are any channels, filters, UDDI references or Web services deployed to the gateway, then remove them.

**Note:** If you have just installed a WebSphere Application Server fix pack or upgrade program, then the upgraded gateway is empty.

2) In the gateway user interface navigation pane, click the following link:

**Gateway**

- [Restore](#)

The **Restore Gateway Configuration** screen is displayed.

3) In the **Location** field, type the path to your configuration file.

**Note:**

- The path to your configuration file must point to a local drive on the machine on which the gateway is currently running.
- The saved file is not specific to a given version of the gateway.
- If you are restoring a file that was backed up using the **Save Gateway Configuration -> Private** option, then the backup contains deployment details (for example URIs) that are specific to a given gateway instance. If you restore this type of backup file to a different instance (for example a gateway on a host with a different network identity), then the restore will succeed but you will have to use the gateway user interface to correct any resultant errors in the deployment details.

4) Click **OK**

Your gateway configuration is restored to the gateway that you are currently administering.

- c. **(Optional)** If you have just restored a configuration after applying an upgrade or fix pack, and any of the deployed Web services were previously published to UDDI registries, use the Listing and managing gateway-deployed Web services option (and the notes you made in step 2 above) to re-publish these Web services to UDDI.

## Backing up and restoring UDDI publication links

When you save a gateway configuration, the deployment details through which a deployed Web service has been published to UDDI by the gateway are not saved. The reasons for not backing up these details are as follows:

- When a gateway is replaced through a fix pack or upgrade program, or otherwise removed without first using the gateway administrative console to un-publish each Web service from UDDI, then the associated entries in UDDI registries are not removed.
- When a gateway is restored, if the UDDI publication details were also restored, then the Web services would be re-published to UDDI by the gateway.
- When a gateway publishes the same service twice to a UDDI registry, the registry does not overwrite the initial publication. It creates a second copy of the TModel and Service definition.
- When a gateway removes a service from a UDDI registry, it only removes the last one published. So it can leave behind defunct TModels and Service Definitions.

So, to preserve the integrity of the UDDI registries, restoring the gateway does not automatically re-publish any Web services to UDDI. But if you know that a service has been un-published from UDDI (for example because you removed it manually before backing up the gateway configuration) then you can safely re-publish it to UDDI after the configuration has been restored.

## Creating and updating a gateway cluster

In order to improve performance and availability, you configure WebSphere Application Server so that a set of application servers acts as a cluster. In a cluster, several application servers are configured identically, and a plug-in to the front-end HTTP server shares incoming HTTP requests between the available application servers in the cluster. Application servers in a cluster can each be hosted on a separate machine, or all hosted on one or two machines. For a more detailed overview of options for clustering in WebSphere Application Server Network Deployment, see *Setting up a multinode environment*.

To create a *gateway* cluster, you manually install an identically-configured gateway on every application server in an existing cluster.

To create and update a gateway cluster, complete the following steps:

### Steps for this task

1. Create a cluster of application servers. For example, by completing the following substeps:

**Note:** These substeps illustrate how to use the WebSphere Application Server Network Deployment deployment manager and administrative console to create a two-server cluster within a cell. In this example, each clustered application server is on a separate node of the cell. However there are other valid configurations, and other ways of creating a cluster. For more detailed information see the InfoCenter for WebSphere Application Server Edge components(<http://www-3.ibm.com/software/webservers/appserv/ecinfocenter.html>), which contains complete documentation for the Caching Proxy and the Load Balancer in the following PDF online books: *WebSphere Application Server Concepts, Planning, and Installation for Edge Components*; the *WebSphere Application Server Caching Proxy Administration Guide*, and the *WebSphere Application Server Programming Guide for Edge Components*.

- a. Install WebSphere Application Server on machines X and Y.
- b. Install WebSphere Application Server Network Deployment on machine X.
- c. Stop all application servers on both machines.
- d. Add machines X and Y as nodes in a deployment manager cell by entering the following command from a command line *on each machine*:

```
WAS_HOME/bin/addnode dm_machine_hostname dm_port -includeapps
```

where

- *WAS\_HOME* is the root directory of your target application server (by default WebSphere/AppServer).
- *dm\_machine\_hostname* is the host name for the machine on which the deployment manager for this cell is running (in this case machine X).
- *dm\_port* is the port on which the deployment manager is listening (by default 8879).

For example (Windows systems):

```
C:\Progra~1\WebSphere\AppServer\bin\addnode xhost 8879 -includeapps
```

- e. Use the WebSphere Application Server Network Deployment administrative console to create new, clustered application servers on the two machines:
  - 1) Select **Servers** -> **Clusters** -> **New**.
  - 2) Enter a name for the cluster.

- 3) Clear the **Prefer local enabled** check box.
- 4) Select the **Do not include an existing server in this cluster** check box.
- 5) Create a new clustered application server on machine X (enter a name for the application server; select the node for this application server (machine X); clear the **generate unique HTTP Ports** check box; click **Apply**).
- 6) Repeat the previous step to create a new clustered application server on machine Y.
- 7) Click **Next**.
- 8) Click **Finish**.

The new cluster is created.

2. Install a gateway on every application server in the cluster.
3. Select and configure the gateway that is to be the source configuration for every gateway in the cluster.

**Note:**

- The gateway that is chosen as the source for the cluster configuration need not be part of the cluster.
  - The channels deployed to this gateway must have their End Point addresses configured for clustering. See Deploying channels to the Web services gateway.
4. Use the back up and restore options to save the source configuration, then restore it to every gateway in the cluster.
  5. Apply all subsequent gateway updates only to the source configuration gateway, then use the back up and restore options to restore the updated configuration to every gateway in the cluster.

**Note:** This is particularly important with regard to Web services that have been published to UDDI by the gateway. When you save a gateway configuration, the record of which Web services have been gateway-published to UDDI is not saved (for reason that are explained in Backing up and restoring UDDI publication links). So if you change the gateway that you use as the source for cluster updates, you lose the record of which Web services have been gateway-published to UDDI.

---

## Administering the Web services gateway

### Before you begin


**5.0.2** If you have just applied a WebSphere Application Server fix pack or upgrade, your gateway configuration is currently empty. If you want to restore a saved configuration, you must do so before you add any new channels, filters, UDDI references or Web services.

To administer the Web services gateway, complete the following steps:

### Steps for this task

1. Start the WebSphere Application Server Administrative Server.
2. Open the following Web page: `http://host:port/wsgw/admin/index.html` where *host* and *port* are the host name and port number that your HTTP server is listening on. For example `localhost:8080` or `localhost:9080`.

The main administration page is displayed:



The order of the elements on this page is significant, for the following reasons:

- If you change the namespace URI or WSDL URI (using the **Configure Gateway** option), you break the link back to the gateway for every Web service that you have already deployed. So you must set these URIs before you deploy any Web services.
- When you configure a Web service, you choose the following entities:
  - The channels on which it is available.
  - The filters (if any) which apply to it.
  - The UDDI references (if any) to which it is deployed.

Each of these choices is made from a list of resources which have already been deployed to the Web services gateway. So you might want to deploy your channels, filters and UDDI references to the gateway before you configure the Web services that use them.

3. For more information on how to configure each element of the Web services gateway, see the following topics:
  - Setting the namespace URI and WSDL URI for the Web services gateway
  - Working with channels
  - Working with filters
  - Working with UDDI references
  - Working with Web services

**Note:**

- You configure each of the above elements of the gateway by filling in fields in a panel.
- In all of the gateway panels, fields marks with asterisks are required.

- After you deploy a channel, filter, or UDDI reference you should refresh all other open browser windows to ensure that up-to-date lists are displayed.

## Setting the namespace URI and WSDL URI for the Web services gateway

### Before you begin

Initial values for the namespace URI and WSDL URI are automatically configured when you install the Web services gateway.

When you deploy a Web service to the Web services gateway, these two URIs are used as follows:

- The **Namespace URI for services** is used as the namespace for the gateway services in exported WSDL documents.
- The **WSDL URI for exported definitions** is used to generate the URL in import statements within exported WSDL documents.

**Note:** When you change these URIs, you break the link back to the Web services gateway for every Web service that you have already deployed. So you must set these URIs before you deploy any Web services to the Web services gateway.

To set the namespace URI and WSDL URI for the Web services gateway, complete the following steps:

### Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **Gateway**

- Configure

The gateway configuration form is displayed:

The screenshot shows the IBM Web Services Gateway administrative interface. On the left is a navigation pane with the following items: Gateway (with sub-item 'Configure'), Services (with sub-items 'List', 'Deploy', 'Remove'), and Channels (with sub-items 'List', 'Deploy'). The main content area is titled 'Configure Gateway' and contains two text input fields. The first field is labeled 'Namespace URI for services' and contains the value 'urn:ibmwsgw'. The second field is labeled 'WSDL URI for exported definitions' and contains the value 'http://PHJ2:9080/wsgw'. Below these fields is an 'Apply Changes' button.

3. In the **Namespace URI for services** field, type the new name.



There is no fixed syntax for the namespace URI, but whatever name you choose is likely to be more effective if it observes the following guidelines:

- It begins with "urn:"  
See the guidance on Internet standards for the syntax of Uniform Resource Names (URNs) at <http://www.ietf.org/rfc/rfc2141.txt>(<http://www.ietf.org/rfc/rfc2141.txt>).
- It is globally unique.
- It reflects your company name.

4. In the **WSDL URI for exported definitions** field, type the new name.

The initial value is the gateway's "best guess" at the right value, but you will probably want to overwrite it with a new value. For example it might guess a local URI such as `http://hlswrth:9080/wsgw`, and because you are giving the WSDL to people in other companies you modify this to `http://hlswrth.your_company.com/wsgw`. Note that only the *host* and *port* parts of the initial value are modified, and that this URI should always start `http://` and end `/wsgw`.

5. Click **Apply Changes**.

What to do next

**5.0.2** **Note:** You also use the gateway configuration form to enable proxy authentication for the gateway.

## Working with channels

Before you begin

Before you can work with a channel, you must install the channel application in WebSphere Application Server as described in the penultimate step of Installing the gateway into a deployment manager cell and Installing the gateway into a stand-alone application server.

Two versions of each type of channel are supplied so that, for each channel type, you can set up separate channels for inbound and outbound requests. For more information see Channels - entry points to the Web services gateway.

From the navigation pane of the Web services gateway administrative user interface, you can choose the following actions for **Channels**:

- **List** to list the deployed channels, and modify their deployment details.
- **Deploy** to deploy a channel.
- **Remove** to remove channels.

### Channels - entry points to the Web services gateway

Channels form entry points to the Web services gateway and carry requests and responses between Web services and the Web services gateway. A request to the Web services gateway arrives through a channel, is translated into a (WSIF) message, then passed through any filters that are registered for the requested service, and finally sent on to the service implementation. Responses follow the same path in reverse.

Before you can use a channel, you must install the channel application in WebSphere Application Server then deploy the channel to the Web services gateway.

**Note:** A deployed channel is not used until you deploy a Web service that uses the channel.

Two versions of each type of channel are supplied with the gateway. This is so that, for each channel type, you can set up separate channels for inbound and outbound requests. This provides a simple mechanism for giving different access rights to users from outside your organisation from the rights you give to users within your organisation:

- To ensure that users outside your organisation can only access those internal services that you choose to publish externally, you deploy those services on the inbound channel.
- To give users inside your organisation access to the full range of internal and external services, you deploy those services on the outbound channel.

### **Listing and managing gateway-deployed channels**

Use this task to list the channels that are deployed to the Web services gateway, and modify their deployment details.

To list the channels that are currently deployed to the Web services gateway, and view and modify their deployment details, complete the following steps:

#### Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **Channels**

- List

The main pane is updated with a list of all the channels that are deployed to the Web services gateway.

3. Click the name of a channel in the list. A form is displayed through which you can view and modify the current deployment details for this channel.
4. Modify the following deployment details:

#### **Home Location**

Type the name of the new home for this channel.

#### **End Point Address**

Type the new address on which the channel is to listen.

#### **5.0.2**

**Note:** For a channel running on an application server that is *not* part of a cluster, this is the address for accessing the channel through the associated HTTP server. For a channel running on an application server that *is* part of a cluster, this address is in two parts (separated by a comma):

- The address for accessing the channel through the associated HTTP server.
  - The address for accessing the channel through the particular application server on which the channel application is running.
5. **(Optional)** If this channel is intended to be used to receive asynchronous reply messages, type appropriate values in the following two fields. Otherwise leave them blank.

**Note:** If the channel supports asynchronous messaging, then the deployment details documentation for the channel should indicate what values to enter in these fields.

#### Async Reply Context Name

#### Async Reply Context Value

6. To start this channel, enable the **YES** radio button. To stop this channel, enable the **NO** radio button.
7. Click **Apply changes**.

#### Results

If the processing completes successfully, the list of deployed channels is redisplayed. Otherwise, an error message is displayed.

## Deploying channels to the Web services gateway

### Before you begin

Before you can deploy a channel, you must install the channel application in WebSphere Application Server as described in the penultimate step of Installing the gateway into a deployment manager cell and Installing the gateway into a stand-alone application server.

If you want to deploy the channels supplied with the Web services gateway, their deployment details are listed in Web services gateway - Channel deployment details.

To deploy a channel, complete the following steps:

#### Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **Channels**

- Deploy

A form is displayed for you to specify the deployment details.

3. Type the following channel deployment details:

#### **Channel Name**

Type the name by which the channel will be known within the Web services gateway (and by which it will be listed using the **Channels > List** option). This name must be unique within the gateway.

#### **Home Location**

Type the name of the home for this channel.

#### **End Point Address**

Type the address on which the channel is to listen.

### **5.0.2**

**Note:** For a channel running on an application server that is *not* part of a cluster, this is the address for accessing the channel through the associated HTTP server. For a channel running on an application server that *is* part of a cluster, this address is in two parts (separated by a comma):

- The address for accessing the channel through the associated HTTP server.
  - The address for accessing the channel through the particular application server on which the channel application is running.
4. **(Optional)** If this channel is intended to be used to receive asynchronous reply messages, type appropriate values in the following two fields. Otherwise leave them blank.

**Note:** If the channel you are deploying supports asynchronous messaging, then the deployment details documentation for the channel should indicate what values to enter in these fields.

**Async Reply Context Name**

**Async Reply Context Value**

5. Click **OK**.

### Results

If the processing completes successfully, the list of deployed channels is updated to include the new channel. Otherwise, an error message is displayed.

### What to do next

**Note:** The deployed channel will not be used until you deploy a Web service that uses the channel.

## **Web services gateway - Channel deployment details**

The deployment details for the channels supplied with the Web services gateway are listed below:

- Apache SOAP channel 1
  - **Channel Name:** ApacheSOAPChannel1
  - **Home Location:** ApacheSOAPChannel1Bean
  - **5.0 +**

### **5.0.1 +**

**End Point Address:** `http://domain_name/wsgwsoap1`

where *domain\_name* is the domain name of the HTTP server through which HTTP requests are passed to this application server.

- **5.0.2 End Point Address:**
  - **Application server not within a cluster:** `http://domain_name/wsgwsoap1`
  - **Application server within a cluster:** `http://domain_name/wsgwsoap1`,  
`http://host_name:port_number/wsgwsoap1`  
where *host\_name* and *port\_number* are the host name and port number for this application server.
  - **Async Reply Context Name:** Leave blank. This function is not supported by this channel.
  - **Async Reply Context Value:** Leave blank. This function is not supported by this channel.
- Apache SOAP channel 2
  - **Channel Name:** ApacheSOAPChannel2

- **Home Location:** ApacheSOAPChannel2Bean

- **5.0+**

- **5.0.1+**

**End Point Address:** `http://domain_name/wsgwsoap2`

where *domain\_name* is the domain name of the HTTP server through which HTTP requests are passed to this application server.

- **5.0.2 End Point Address:**

- **Application server not within a cluster:** `http://domain_name/wsgwsoap2`

- **Application server within a cluster:** `http://domain_name/wsgwsoap2`,  
`http://host_name:port_number/wsgwsoap2`

where *host\_name* and *port\_number* are the host name and port number for this application server.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.

- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

- **5.0.2** SOAP/HTTP Channel 1

- **Channel Name:** SOAPHTTPChannel1

- **Home Location:** SOAPHTTPChannel1Bean

- **End Point Address:**

- **Application server not within a cluster:**

`http://domain_name/wsgwsoaphttp1`

- **Application server within a cluster:** `http://domain_name/wsgwsoaphttp1`,  
`http://host_name:port_number/wsgwsoaphttp1`

where *host\_name* and *port\_number* are the host name and port number for this application server.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.

- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

- **5.0.2** SOAP/HTTP Channel 2

- **Channel Name:** SOAPHTTPChannel2

- **Home Location:** SOAPHTTPChannel2Bean

- **End Point Address:**

- **Application server not within a cluster:**

`http://domain_name/wsgwsoaphttp2`

- **Application server within a cluster:** `http://domain_name/wsgwsoaphttp2`,  
`http://host_name:port_number/wsgwsoaphttp2`

where *host\_name* and *port\_number* are the host name and port number for this application server.

- **Async Reply Context Name:** Leave blank. This function is not supported by this channel.

- **Async Reply Context Value:** Leave blank. This function is not supported by this channel.

## Removing channels from the Web services gateway

To remove a channel, complete the following steps:

### Steps for this task

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **Channels**

- Remove

The main pane is updated with a list of all the channels that are deployed to the Web services gateway. Alongside each entry in the list is a check box, and information on the number of Web services that currently use the channel.

3. **(Optional)** Click the name of a channel in the list.

A form is displayed through which you can view the current deployment details for this channel, including a list of the Web services that currently use the channel.

4. Select the check box for every channel that you want to remove.

**Note:** When you remove a channel that is currently used by one or more Web services, the gateway removes the channel from the channel list for each associated Web service.

5. Click **OK**.

### Results

If the processing completes successfully, the list of deployed channels is updated. Otherwise, an error message is displayed.

## Working with filters

### Before you begin

Before you can work with a filter, you must install the filter application in WebSphere Application Server as described in the penultimate step of Installing the gateway into a deployment manager cell and Installing the gateway into a stand-alone application server.

From the navigation pane of the Web services gateway administrative user interface, you can choose the following actions for **Filters**:

- **List** to list the deployed filters, and modify their deployment details.
- **Deploy** to deploy a filter.
- **Remove** to remove filters.

You can also write your own filters for the Web services gateway.

### **Filters - service interceptors for the Web services gateway**

Filters are used to intercept service invocations which come into the Web services gateway, and responses which leave it. Filters can perform a wide range of tasks, from logging messages, to transforming their content, to terminating an incoming request. Filters are deployed to the Web services gateway as described in Deploying filters to the Web services gateway, then registered for use with individual Web services as described in "Working with Web services".

## Listing and managing gateway-deployed filters

Use this task to list the filters that are deployed to the Web services gateway, and modify their deployment details.

To list the filters that are currently deployed to the Web services gateway, and view and modify their deployment details, complete the following steps:

### Steps for this task

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **Filters**

- [List](#)

The main pane is updated with a list of all the filters that are deployed to the Web services gateway.

3. Click the name of a filter in the list. A form is displayed through which you can view and modify the current deployment details for this filter.
4. Modify the following deployment detail:

#### **Home Location**

Type the name of the new home for this filter.

5. Click **Apply changes**.

### Results

If the processing completes successfully, the list of deployed filters is redisplayed. Otherwise, an error message is displayed.

## Deploying filters to the Web services gateway

### Before you begin

Use this task to deploy a filter to the Web services gateway.

**Note:** The deployed filter will not be used until you deploy a Web service that uses the filter.

Before you can deploy a filter, you must install the filter application in WebSphere Application Server as described in the penultimate step of Installing the gateway into a deployment manager cell and Installing the gateway into a stand-alone application server.

**Note:** You can deploy multiple instances of a filter by entering different filter names.

To deploy a filter, complete the following steps:

### Steps for this task

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **Filters**

- [Deploy](#)

A form is displayed for you to specify the deployment details.

3. Type the following filter deployment details:

**Filter Name**

Type the name by which the filter will be known within the Web services gateway (and by which it will be listed using the **Filters > List** option). This name must be unique within the Web services gateway.

**Home Location**

Type the name of the home for this filter.

4. Click **OK**.

Results

If the processing completes successfully, the list of deployed filters is updated to include the new filter. Otherwise, an error message is displayed.

**Removing filters from the Web services gateway**

To remove a filter, complete the following steps:

Steps for this task

1. Display the Web Services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Filters**

- Remove

The main pane is updated with a list of all the filters that are deployed to the Web services gateway. Alongside each entry in the list is a check box, and information on the number of Web services that currently use the filter.

3. **(Optional)** Click the name of a filter in the list.

A form is displayed through which you can view the current deployment details for this filter, including a list of the Web services that currently use the filter.

4. Select the check box for every filter that you want to remove.

**Note:** When you remove a filter that is currently used by one or more Web services, the gateway removes the filter from the filter lists for each associated Web service.

5. Click **OK**.

Results

If the processing completes successfully, the list of deployed filters is updated. Otherwise, an error message is displayed.

**Working with UDDI references**

A UDDI reference is a pointer to a UDDI registry. This may be a private UDDI registry such as the (IBM WebSphere UDDI Registry), or a public UDDI registry.

In the UDDI model, Web services are owned by "businesses", and "businesses" are owned by "Authorized Names". Each UDDI reference gives access to the Web services that are owned by a single "Authorized Name" in single UDDI registry.

From the navigation pane of the Web services gateway administrative user interface, you can choose the following actions for **UDDI References**:

- **List** to list the deployed UDDI references, and modify their deployment details.
- **Deploy** to deploy a UDDI reference.



- **Remove** to remove UDDI references.

For more information about how the gateway works with UDDI registries, see *UDDI registries - Web service directories that integrate with the Web services gateway* and *Publishing a Web service to a UDDI registry for deployment to the gateway*. For more general information about UDDI and UDDI registries, see the UDDI community at [uddi.org](http://uddi.org)(<http://uddi.org>).

## **UDDI registries - Web service directories that integrate with the Web services gateway**

### **UDDI**

The Universal Description, Discovery and Integration (UDDI) specification defines a way to publish and discover information about Web services.

In this specification:

- Each Web service is owned by one "business", and each "business" (and the Web services it owns) is maintained by one "Authorized Name".
- One "Authorized Name" can own many "businesses", and one "business" can own many Web services.

The UDDI specification also associates Web services with "Technical models". These are generic categories that allow a UDDI registry user to search for a *type* of service, rather than needing to know the access details for a specific service.

For more general information about UDDI, see the UDDI community at [uddi.org](http://uddi.org)(<http://uddi.org>).

### **UDDI registries**

UDDI registries use the UDDI specification to publish directory listings of Web services. There are Universal Business Registries (sometimes referred to as 'public UDDI registries') hosted worldwide, including one hosted by IBM. Enterprises can also host their own internal registries behind their firewalls (sometimes referred to as 'private UDDI registries') to better manage their internal implementation of Web services. The (WebSphere UDDI Registry) is an example of a private UDDI registry.

### **How the gateway interacts with UDDI registries**

The gateway interacts with UDDI registries in two ways:

- When you deploy a Web service to the gateway, you specify the location of the "internal" WSDL file that describes the Web service to be deployed. This WSDL file can be located through a UDDI registry.
- For any gateway-deployed Web service, you can tell the gateway to create entries for the Web service in one or more UDDI registries.

To enable your gateway to interact with a UDDI registry, you create one or more gateway pointers to the registry. The gateway refers to these pointers as "UDDI references", and you create them as described in *Deploying UDDI references to the Web services gateway*. Each UDDI reference includes the following parameters:

- The access points for the UDDI registry (the **Inquiry URL** and the **Publish URL**).
- The "Authorized Name" (the **User ID** and **Password**) for the owner of one or more "businesses" in the UDDI registry.

You get the "Authorized Name" from the target UDDI registry. For more information see Publishing a Web service to a UDDI registry for deployment to the gateway.

**Note:** A given UDDI reference can only access the Web services that are owned by the "businesses" that are in turn owned by a single "Authorized Name". So if you need access to two Web services in the same registry, and each service is owned by a different "Authorized Name", then you need to create two UDDI references.

When you deploy a Web service, and you specify that the "internal" WSDL file is located through a UDDI registry, you enter the following two parameters:

- The UDDI reference that can access this service.
- The "service key" that the UDDI registry has assigned to this service.

You get the "service key" from the target UDDI registry. For more information see Publishing a Web service to a UDDI registry for deployment to the gateway

When you tell the gateway to create entries for a deployed Web service in one or more UDDI registries, you enter the following two parameters:

- The UDDI references (one for each registry) that can access the UDDI business category under which you want to publish this service.
- The "business key" that identifies the UDDI business category.

You get the "business key" from the target UDDI registry. For more information see Publishing a Web service to a UDDI registry for deployment to the gateway

**Note:** Because the gateway only interacts with UDDI registries at the level of specific Web services, the gateway does not make use of UDDI "technical models".

### **Listing and managing gateway-deployed UDDI references**

Use this task to list the UDDI references that are deployed to the Web services gateway, and modify their deployment details.

To list the UDDI references that are currently deployed to the Web services gateway, and view and modify their deployment details, complete the following steps:

#### Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **UDDI References**

- List

The main pane is updated with a list of all the UDDI references that are deployed to the Web services gateway.

3. Click the name of a UDDI reference in the list. A form is displayed through which you can view and modify the current deployment details for this UDDI reference.
4. Modify the following deployment details:

#### **Inquiry URL**

Type the new URL that provides access to this registry for the SOAP inquiry API.

**Publish URL**

Type the new URL that provides access to this registry for the SOAP publish API.

**User Name**

Type the new user ID for the "Authorized Name" that has update access to this registry.

**Password**

Type the password for this new user ID.

**Confirm Password**

Type again the password for this new user ID.

**Note:**

The values you enter here for **User Name** and **Password** must match those of the owner of the corresponding business in UDDI. You can see the owning user ID in UDDI by looking at the business details under the "Authorized Name" field.

If the values you enter here do not match the "Authorized Name" values for the business that owns the service, then the service will not be published or found.

If the business that owns the service has more than one "Authorized Name", you might want to set up multiple UDDI references (each with a different user ID) to the same UDDI registry .

5. Click **Apply changes**.

Results

If the processing completes successfully, the list of deployed UDDI references is redisplayed. Otherwise, an error message is displayed.

**Deploying UDDI references to the Web services gateway**Before you begin

**Note:** The deployed UDDI reference will not be used until you deploy a Web service that uses the UDDI reference.

To deploy a UDDI reference, complete the following steps:

Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

**UDDI References**

- Deploy

A form is displayed for you to specify the deployment details.

3. Type the following UDDI reference deployment details:

**Reference Name**

Type the name by which the UDDI reference will be known within the Web services gateway (and by which it will be listed using the **UDDI References > List** option). This name must be unique within the gateway - and note that you might need more than one UDDI reference

for a given UDDI registry (for more information see UDDI registries - Web service directories that integrate with the Web services gateway).

**Inquiry URL**

Type the URL that provides access to this registry for the SOAP inquiry API.

**Publish URL**

Type the URL that provides access to this registry for the SOAP publish API.

**User Name**

Type the user ID for an "Authorized Name" that has update access to this registry.

**Password**

Type the password for this user ID.

**Confirm Password**

Type again the password for this user ID.

**Note:**

The values you enter here for **User Name** and **Password** must match those of the "Authorized Name" in the UDDI registry. You can see the owning user ID in UDDI by looking at the business details under the "Authorized Name" field.

4. Click **OK**.

Results

If the processing completes successfully, the list of deployed UDDI references is updated to include the new UDDI reference. Otherwise, an error message is displayed.

**Removing UDDI references from the Web services gateway**

To remove a UDDI reference, complete the following steps:

Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

**UDDI References**

- Remove

The main pane is updated with a list of all the UDDI references that are deployed to the Web services gateway. Alongside each entry in the list is a check box, and information on the number of Web services that currently use the UDDI reference.

3. **(Optional)** Click the name of a UDDI reference in the list.

A form is displayed through which you can view the current deployment details for this UDDI reference, including a list of the Web services that currently use the UDDI reference.

4. Select the check box for every UDDI reference that you want to remove.

**Note:** When you remove a UDDI reference that is currently used by one or more Web services, the gateway removes the UDDI reference from the UDDI reference list for each associated Web service.

5. Click **OK**.

## Results

If the processing completes successfully, the list of deployed UDDI references is updated. Otherwise, an error message is displayed.

## **Working with Web services**

### Before you begin

If you change the Namespace URI, you break the link back to the Web services gateway for every Web service that you have already deployed. So you must set the Namespace URI before you deploy any Web services.

When you configure a Web service, you choose the following resources:

- The channels on which the service is available.
- The filters (if any) that apply to the service.
- The UDDI registries (if any) in which entries for the service are created.

Each of these choices is made from a list of resources that have already been deployed to the Web services gateway. So you must deploy your channels, filters and references to UDDI registries to the Web services gateway before you deploy the Web services that use those resources.

From the navigation pane of the Web services gateway administrative user interface, you can choose the following actions for **Services**:

- **List** to list the deployed Web services, and modify their deployment details.
- **Deploy** to deploy a Web service.
- **Remove** to remove Web services.

### **Listing and managing gateway-deployed Web services**

Use this task to list the Web services that are deployed to the Web services gateway, and modify their deployment details.

### Before you begin

There is no point in deploying multiple target services to the same gateway service unless you have a filter implementation that is capable of selecting the required target service.

For information on how to write pluggable filters that can select the target service, see [Writing a filter for the Web services gateway](#) and [Using a filter to select a target service and port](#).

To list the Web services that are currently deployed to the Web services gateway, and view and modify their deployment details (including adding or removing multiple target services) complete the following steps:

### Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **Services**

- List

The main pane is updated with a list of all the Web services that are deployed to the Web services gateway.

3. Click the name of a Web service in the list. A form is displayed through which you can view and modify the current deployment details for this Web service, and add or remove multiple target services.
4. At the level of the gateway service itself (in the **Gateway Service Properties** section) you can change the following settings. When you have finished making changes, click **Apply Changes**.

- a. **Authorization Policy - Control access to this service.** Use this check box to enable or disable operation-level authorization for this gateway service.

- b. **Audit Policy - Log requests to this service.**

The Audit policy indicates whether the MessageWarehouse object, if present, should be used to log requests and responses for this service. If you have a Message Warehouse implementation, use this check box to enable or disable logging of requests and responses for this Web service.

For more information see Capturing Web service invocation information and The MessageWarehouse interface.

- c. In this release of the gateway, the **Annotation URL** field is not used.

- d. If you want to publish the service to one or more UDDI registries (selected in the **UDDI References** section below), enter the UDDI business key in the field provided under **UDDI Publication Properties**. This key identifies the business category under which you want your service to appear in UDDI. To get a list of valid business keys, look up businesses in a UDDI registry. This is an example of a UDDI business key: 08A536DC-3482-4E18-BFEC-2E2A23630526. For more information about UDDI business keys see Publishing a Web service to a UDDI registry for deployment to the gateway.

- e. If you want to enable or disable **Service Security**, select or clear the **Enable service security** check box. If you choose to enable service security, then you must use the **Edit service security configuration** option to configure secure communication for this gateway service between the service requester (the client) and the gateway. For more information see Editing the service security configuration.

5. In the **Target Services** section you can add, modify or remove services from a list of target services for this single gateway service. Every service on this list provides exactly the same service, and they are presented by the gateway to the service requesters as a single gateway service. To add a new target service, complete the following steps:

- a. **WSDL Location.** Type the location of the "internal" WSDL file that describes the Web service to be deployed. The WSDL file is either located at a URL, or through a UDDI registry.

If the WSDL location is a URL, type the URL (if the binding and service definition for this Web service are held in separate WSDL files, then type the URL of the WSDL file that defines the binding).

If the WSDL is located through a UDDI registry, type *uddiReference,serviceKey* where

- *uddiReference* is the reference name by which a currently-deployed UDDI reference is known within the gateway (and by which it is listed using the **UDDI References -> List** option)
- *serviceKey* is the service key that the UDDI registry has assigned to the service. This is an example of a UDDI service key: 34280367-0ECF-46CE-

B804-14C21D6D0FB1. For more information about UDDI service keys see Publishing a Web service to a UDDI registry for deployment to the gateway.

**Note:**

- When the Web services gateway deploys the Web service, it generates a matching "external" WSDL file that it makes available to gateway users. This "external" WSDL file also describes the service, but is located at a new URL and is generated and maintained by the Web services gateway itself.
  - If the Web service is also published to one or more UDDI registries, then the "internal" WSDL file is required to remove the service from the gateway.
- b. **Location Type.** Select the type of location (either URL or UDDI) where the "internal" WSDL file is held.
  - c. **Target Service Name.** If the Web service WSDL contains more than one service, or the WSDL is located through a UDDI registry, type the target service's name from the target service WSDL.
  - d. **Target Service Namespace.** If the Web service WSDL contains more than one service, or the WSDL is located through a UDDI registry, type the namespace of the target service's name from the target service WSDL.
  - e. **Target Service Identity Information.** Type the identity by which the target service is known within the Web services gateway. This identity need not be unique.  
**Note:** If you are mapping multiple target services, and also writing a filter to select a target service, you might use the Target Service Identity Information to select a particular target service from the set.
  - f. Click **add**.  
The target service is added to a list of target services.

**5.0.2** To modify the deployment details for an existing target service, click on the name of the service in the list of target services. A form is displayed, containing the same fields that you filled in when you added the service, and also the following additional fields:

- **Started.** Enable this check box option if you want the gateway to offer this Web service to service requesters.
- **Enable target service security.** Enable this check box if you want the gateway to use secure communication between the gateway and the target service.
- **Edit target service security configuration.** If you choose to enable target service security, then you must use this option to configure secure communication between the gateway and the target service. For more information see Editing the target service security configuration.

When you have finished making changes, click **Apply Changes**.

6. In the **Channels** section, you can add or remove channels from the list of deployed channels through which this service is available.
7. In the **Request Filters** section, you can add or remove filters from the list of deployed filters that are applied to the request.

**Note:** The filters are executed in the order shown. To add a filter into the list at a particular position, use the **at position** menu.

8. In the **Response Filters** section, you can add or remove filters from the list of deployed filters that are applied to the response.  
**Note:** The filters are executed in the order shown. To add a filter into the list at a particular position, use the **at position** menu.
9. In the **UDDI References** section, you can add or remove UDDI references from the list of deployed UDDI references to UDDI registries in which this service is published. If you select one or more UDDI references in this step, you must also enter a UDDI business key in the field provided under **UDDI Publication Properties** as described above. For more information about how the gateway works with UDDI registries, see UDDI registries - Web service directories that integrate with the Web services gateway.
10. In the **Exported WSDL definitions** section there are two pairs of WSDL links. Both pairs link to (a) the external WSDL implementation definition, and (b) the external WSDL interface definition.
  - To view details of the associated external WSDL for the service, use the first pair.
  - To return the WSDL for use by application programs that need the WSDL definitions for the service, use the second pair.

If there is an error generating the WSDL then a blank page is returned.

**Note:** To help your service users locate the WSDL documents for services that are deployed to the Web services gateway, the gateway also supports the WS-Inspection specification(<http://www.ibm.com/developerworks/webservices/library/ws-wsilspec.html>). To open a WS-Inspection document which contains references to the WSDL documents for all of the gateway-deployed services, you issue an HTTP GET against

```
http://host:port/wsgw/wsinspection.wsil
```

where *host* and *port* are the host name and port number that your HTTP server is listening on.

## Deploying Web services to the Web services gateway

Before you begin

### 5.0+

#### 5.0.1+

Before you deploy a Web service to the gateway, deploy the resources (channels, filters, and UDDI references) that the Web service uses.

**5.0.2** Before you deploy a Web service to the gateway, deploy the resources (channels, filters, UDDI references and security bindings) that the Web service uses.

To deploy a Web service, complete the following steps:

Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Services**

- Deploy



A form is displayed for you to specify the deployment details.

3. Type the following Web service deployment details:
  - a. **Gateway Service Name.** Type the name by which the Web service will be known within the gateway, and by which it will be listed using the Services > List option. This name must be unique within the gateway and must not contain any spaces.

**Note:** If you have several implementations of the same Web service, you can map them all to the same deployed gateway service. You use the **Deploy** option (this option) to deploy just one instance of a given gateway service. To add more target services to a deployed gateway service, you use the Services > List option.
  - b. Choose between the two types of **Message part representation**. For more information see Data type representation - choosing between Generic classes and Deployed Java classes.

#### Generic classes

You can select this option in all cases except the following cases:

- You cannot select this option if your service has EJB or Java bindings that are used at runtime.
- You cannot select this option if your service uses Vectors, Enumerations, Hashtables or Maps.

#### Deployed Java classes

You must select this option if the target services for the gateway service use Vectors, Enumerations, Hashtables or Maps, or contain EJB or Java bindings.

**Note:** For Web services deployed with Java bindings (or EJB bindings even where the Web service is on a different server) you must also make sure that the specific Java classes that have been generated for the Web service are available to the gateway. For more information see Deploying Web services with Java bindings.

You can also select this option (as an alternative to selecting "Generic classes"):

- If the target service uses only supported simple and compound types.
- If the target service uses complex types, and the specific Java classes that have been generated for the Web service are available to the gateway (but note that if you select "Generic classes", then you don't need to also deploy the associated Java classes locally).

- c. **(Optional) Authorization Policy - Control access to this service.** If you want to enable operation-level authorization for this Web service, enable this check box.
- d. **(Optional) Audit Policy - Log requests to this service.**

The Audit policy indicates whether the MessageWarehouse object, if present, should be used to log requests and responses for this service. If you have a Message Warehouse implementation, and you want it to log requests and responses for this Web service, enable this check box.

For more information see Capturing Web service invocation information and The MessageWarehouse interface.
- e. In this release of the gateway, the **Annotation URL** field is not used.

- f. Select the deployed resources that the Web service is to use, from the following lists:

**Channels**

Select one or more deployed channels through which this service is to be available.

**Request Filters**

Select zero or more deployed filters to apply to the request.

**Response Filters**

Select zero or more deployed filters to apply to the response.

**UDDI References**

Select zero or more deployed UDDI references (one for each UDDI registry) that can access the UDDI business category under which you want to publish this service. If you select one or more UDDI references in this step, you must also enter the UDDI business key in step 3h below. For more information about how the gateway works with UDDI registries, see UDDI registries - Web service directories that integrate with the Web services gateway.

- g. Specify the **Target Service Properties** for the Web service:

**WSDL Location**

Type the location of the "internal" WSDL file that describes the Web service to be deployed. The WSDL file is either located at a URL, or through a UDDI registry.

If the WSDL location is a URL, type the URL (if the binding and service definition for this Web service are held in separate WSDL files, then type the URL of the WSDL file that defines the binding).

If the WSDL is located through a UDDI registry, type *uddiReference,serviceKey* where

- *uddiReference* is the reference name by which a currently-deployed UDDI reference is known within the gateway (and by which it is listed using the **UDDI References -> List** option)
- *serviceKey* is the service key that the UDDI registry has assigned to the service. This is an example of a UDDI service key: 34280367-0ECF-46CE-B804-14C21D6D0FB1. For more information about UDDI service keys see Publishing a Web service to a UDDI registry for deployment to the gateway.

**Note:**

- When the Web services gateway deploys the Web service, it generates a matching "external" WSDL file that it makes available to gateway users. This "external" WSDL file also describes the service, but is located at a new URL and is generated and maintained by the Web services gateway itself.
- If the Web service is also published to one or more UDDI registries, then the "internal" WSDL file is required to remove the service from the gateway.

**Location Type**

Select the type of location (either URL or UDDI) where the "internal" WSDL file is held.

**Target Service Name**

If the Web service WSDL contains more than one service, or the

WSDL is located through a UDDI registry, type the target service's name from the target service WSDL.

#### **Target Service Namespace**

If the Web service WSDL contains more than one service, or the WSDL is located through a UDDI registry, type the namespace of the target service's name from the target service WSDL.

#### **Target Service Identity Information**

Type the identity by which the Web service is known within the Web services gateway. This identity need not be unique.

**Note:** If you later add more target services to this gateway service, and also write a filter to select a target service, you might use the Target Service Identity Information to select a particular target service from the set.

- h. If you want to publish the service to one or more UDDI registries (selected in a previous step), enter the UDDI business key in the field provided under **UDDI Publication Properties**. This key identifies the business category under which you want your service to appear in UDDI. To get a list of valid business keys, look up businesses in a UDDI registry. This is an example of a UDDI business key: 08A536DC-3482-4E18-BFEC-2E2A23630526. For more information about UDDI business keys see Publishing a Web service to a UDDI registry for deployment to the gateway.

4. Click **OK**.

#### Results

If the processing completes successfully, the list of deployed Web services is updated to include the new Web service. Otherwise, an error message is displayed.

#### What to do next

##### **5.0+**

##### **5.0.1+**

After deployment, use the list deployed Web services option to change the resources (channels, filters and UDDI references) that the Web service uses, or to add multiple target services for this gateway service.

##### **5.0.2**

After deployment, use the list deployed Web services option to change the resources (channels, filters, UDDI references and security bindings) that the Web service uses, or to add multiple target services for this gateway service.

If you enabled the check box 'Authorization Policy - Control access to this service', you must now enable Web service operation-level authorization.

### **Data type representation - choosing between Generic classes and Deployed Java classes**

When you deploy a Web service, the **Message part representation** option allows you to choose between **Generic classes** and **Deployed Java classes**.

As your message passes through the gateway, the message parts are represented as actual Java objects. The data type used for each part is defined as follows:

- It is one of the set of XML schema and SOAP supported simple and compound data types, or
- It is a complex type defined in the WSDL schema section.

**Note:** A complex type is a data type represented by a Java class (such as a user-written class) that is not part of the native Java language.

**Generic classes** and **Deployed Java classes** can both represent simple, compound and complex data types, subject to the following constraints:

- The gateway only supports the simple and compound types that are listed in Web services gateway - Supported types.
- Only **Deployed Java classes** can represent Vectors, Enumerations, Hashtables and Maps (but either **Generic classes** or **Deployed Java classes** can represent Arrays).
- If **Deployed Java classes** are used to represent complex types, then the actual Java classes representing these complex types must be deployed to the application server on which the gateway is running.
- If the target service uses Java or EJB WSDL bindings (that is, if the target service is a Java class deployed on the local application server, or it is an enterprise bean) then **Deployed Java classes** must be used, and the bindings must be made available as described in Deploying Web services with Java bindings.

The gateway's schema parser determines all top-level types that are defined in the WSDL schema section, and generates mappings to **generic classes** for all of these types. This enables the gateway to forward requests (and responses) containing most complex data type parameters (and return values) to a remote destination without requiring the actual Java classes representing these complex types to be deployed to the application server on which the gateway is running. So if your Web service uses complex data types, and there is no other constraint that forces you to use **Deployed Java classes**, then you should select **Generic classes**.

Performance is the same whether you choose to use **Deployed Java classes** or **Generic classes**.

### **Complex data types - mapping namespaces to packages**

If you write your own WSDL file that describes your Web service (rather than use an automated tool such as WebSphere Studio Application Developer) and the service uses complex data types, then to ensure that it can be successfully deployed to the gateway you should follow these guidelines for mapping namespaces to packages.

For working with complex data types, there is no industry-wide standard way of mapping namespaces to packages (in fact the JAX-RPC standard states that the tools themselves must make up their own standard). The standard used by the gateway is as follows:

- Set the namespace of the complex data type to the class's java package name.
- Set the complex data type name to the class's name.

So if the java class you are using for the complex data type is `random.RandomData`, then the namespace of the complex data type is `random`, and the complex data type name is `RandomData`. And if the full package name is `com.ibm.www.random`, then the namespace is `www.ibm.com/random`.

See also the troubleshooting tip about working with Web services that use complex data types.

### Deploying Web services with Java bindings

For Web services deployed with Java bindings (or EJB bindings even where the Web service is on a different server) you must make additional classes available to the gateway.

For **EJB bindings**, make the EJB client JAR file available. If the Web service is deployed on the same server as the gateway, the necessary interfaces and classes are already visible. If not, you should implement one of the following options:

- Copy the EJB client JAR file into the *WAS\_HOME/lib* or *WAS\_HOME/lib/app* directory (where *WAS\_HOME* is the root directory for your installation of IBM WebSphere Application Server).
- Update the application server class path to include the EJB client JAR file.

For **Java bindings**, make the Java classes for the Web service available by implementing one of the following options:

- Copy the JAR or class files that contain the Java classes into the *WAS\_HOME/lib* or *WAS\_HOME/lib/app* directory.
- Update the application server class path to include the JAR file.
- Wrap the Java classes in an enterprise bean and deploy it on the same application server. WebSphere Application Server will then make the classes available to the gateway application.

### Web services gateway - Supported types

When you deploy a Web service, you use the **Message part representation** option to choose between Generic classes and Deployed Java classes. This topic gives reference information about the data types that are supported in each case.

#### Simple types

The following table gives a list of the XML schema (and Java equivalent) simple types that are supported by both **Generic classes** and **Deployed Java classes**:

##### XML schema (and Java equivalent) simple type

xsd:string  
xsd:float  
xsd:double  
xsd:int  
xsd:boolean  
xsd:byte  
xsd:short  
xsd:long  
xsd:decimal  
xsd:QName  
xsd:date  
xsd:timeInstant

#### Compound types

From a Java perspective, compound types are types with constituent elements. These elements are either identified purely by name (for example, a Java class with several member properties) or by ordinal position (for example, a List data structure like Array or Vector).

The following table gives a list of the SOAP (and Java equivalent) compound types that are supported:

| SOAP compound type | java equivalent compound type | Supported by                              |
|--------------------|-------------------------------|---|
| Array              | Java array                    | Generic classes and Deployed Java classes |
| Vector             | java.util.Vector              | Deployed Java classes                     |
| Enumeration        | java.util.Enumeration         | Deployed Java classes                     |
| Map                | java.util.Hashtable           | Deployed Java classes                     |
| Map                | java.util.Map                 | Deployed Java classes                     |

### Complex types

A complex type is a data type represented by a Java class (such as a user-written class) that is not part of the native Java language.

Complex types can include combinations of simple types, compound types, and other complex types. For example, the children of a complex type might be represented by another complex type, or by any of the simple or compound types.

**Generic classes** and **Deployed Java classes** can both represent complex types, subject to the constraints described in Data type representation - choosing between Generic classes and Deployed Java classes.

### Publishing a Web service to a UDDI registry for deployment to the gateway

The gateway interacts with UDDI registries as described in UDDI registries - Web service directories that integrate with the Web services gateway. When you deploy a Web service to the gateway, you enable UDDI interaction by entering a UDDI reference, and (depending upon what you are trying to do) either or both of the following pieces of information:

- The "service key" that the UDDI registry has assigned to this service.
- The "business key" that identifies the UDDI business category under which you want your service to appear in the UDDI registries.

You get these two keys *from the UDDI registry*. To help you understand what UDDI "service keys" and "business keys" are, and where you find them in a UDDI registry, this topic describes how to publish a Web service to a UDDI registry.

#### Note:

- The UDDI publication process described below requires that you specify a "technical model". Technical models are generic categories. They allow a UDDI registry user to search for a *type* of service, rather than needing to know the access details for a specific service. The gateway makes no use of technical model information, because it only interacts with UDDI registries at the level of specific Web services.
- The following task steps include specific navigation instructions. These instructions describe how you publish a Web service to the (IBM WebSphere UDDI Registry). If you are working with a different UDDI registry, then the specific navigation will be different but the underlying principles will be the same.

#### Steps for this task

1. Specify a business:

- a. To get a list of valid business keys, look up businesses in the UDDI registry.  
This is an example of a UDDI business key: 08A536DC-3482-4E18-BFEC-2E2A23630526.
  - b. If you do not find an appropriate existing business in the UDDI registry, then use the **Add a business** option on the **Advanced Publish** section of the Publish pane to add a new one.
2. Add a technical model:
    - a. Select **Add a technical model** on the **Advanced Publish** section of the Publish pane.
    - b. Enter the name as specified for the target namespace of your binding (or interface) WSDL file, then add a description (if required).
    - c. Add a category of Type unspsc and value wsdlSpec (the Key name field can be left blank).
    - d. Add an overview URL specifying the URL for your binding WSDL file, then add a description (if required).  
**Note:** The binding and the service definition for your Web service might be held in separate WSDL files, so be careful to type the URL of the WSDL file that defines the *binding*.
    - e. Click **Publish Technical Model**.
  3. Add a service:
    - a. Select **Show owned entities** on the **Advanced Publish** section of the Publish pane.
    - b. Select **Add a Service** for your business.
    - c. Enter the name as specified for the target service in your WSDL file, then add a description (if required).
    - d. For the **Access point** ensure the correct URL type is selected (for example http for an http access point), then enter the value of the soap:address location (or its equivalent) from your service definition WSDL file (for example http://yourhost:80/SimpleTest/servlet/rpcrouter).
    - e. For the **Technical model** select **Add**, then find the required technical model by entering a suitable prefix and selecting **Find technical models**, then check the selection box for the required technical model and click **Update**.
    - f. Click **Publish Service**.

### Results

The UDDI registry assigns a **service key** to your service, and publishes the service.

### **Removing Web services from the Web services gateway**

To remove a Web service, complete the following steps:

#### Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### **Services**

- Remove

The main pane is updated with a list of all the Web services that are deployed to the Web services gateway. Alongside each entry in the list is a check box.

3. Select the check box for every Web service that you want to remove.

4. Click **OK**.

### Results

If the processing completes successfully, the list of deployed Web services is updated. Otherwise, an error message is displayed.

**Note:** If the Web service that you want to remove is also published to one or more UDDI registries, then the "internal" WSDL file is required to remove the service from the gateway. So if the service is published to UDDI, and the processing does not complete successfully, check that the WSDL file is still available at the location defined for the service in the **Target Services -> WSDL Location**. For more information see Listing and managing gateway-deployed Web services.

---

## Running the Web services gateway samples

The following pre-built samples are available for use with the Web services gateway:

- The standard Stock Quote service sample, that requires an Internet connection.
- The Address Book service sample, that allows the storing and retrieval of names and addresses.

### Results

These samples, and documentation on how to use them, are available through the Web services gateway samples link on the Samples Central (<http://www.ibm.com/websphere/developer/library/samples/AppServer.html>) page of the IBM WebSphere Developer Domain Web site.

### What to do next

If you want to test the gateway taking service definitions from a private UDDI registry such as the (WebSphere UDDI Registry), you should complete the following additional steps:

1. Publish the WSDL for each of these samples to UDDI. (For more information on how to do this, see Publishing a Web service to a UDDI registry for deployment to the gateway and the documentation for your private UDDI registry).
2. Instruct the gateway to locate the service through the UDDI registry, as described in Deploying Web services to the Web services gateway.

---

## Passing SOAP messages with attachments through the Web services gateway

The Web services gateway supports Web services that pass attachments in a MIME message. This support is included in the SOAP/HTTP channel.

Attachments are carried through the various gateway components and passed on to the target service. The content MIME type of each attachment is preserved.

When the Target service is deployed to a JAX-RPC compliant server, the attachments can be accessed on the target service using `javax.activation.DataHandler`.



The WSDL representing a SOAP messages with attachments service must define the attachment parts in the Binding section. The `mime:multipartRelated`, `mime:part` and `mime:content` tags are used to describe the attachment.

For more information, see the following topics:

- SOAP messages with attachments - a definition.
- Writing the WSDL extensions for SOAP messages with attachments.

**Note:** The following scenarios are not supported:

- Using the Apache SOAP channel.
- Using DIME.
- Using the `mime:mimeXml` WSDL tag.
- Nesting a `mime:multipartRelated` inside a `mime:part`.
- Using Arrays or Vectors of DataHandlers, Images, and so forth.

The MIME headers from the incoming message are not preserved for referenced attachments. The outgoing message contains new MIME headers for Content-Type, Content-Id and Content-Transfer-Encoding that are created by (WSIF).

## SOAP messages with attachments - a definition

From an architecture and external specification viewpoint, SOAP Messages with Attachments(<http://www.w3.org/TR/SOAP-attachments>) is an extension to the SOAP 1.1 (<http://www.w3.org/TR/SOAP/>) Recommendation from the World Wide Web Consortium (W3C)(<http://www.w3.org/>).

The W3C SOAP messages with attachments document describes a standard way to associate a SOAP message with one or more attachments in their native format (for example GIF or JPEG) by using a multipart MIME structure for transport. It defines specific usage of the Multipart/Related MIME media type, and rules for the usage of URI references to refer to entities bundled within the MIME package. It thereby outlines a technique for a SOAP 1.1 message to be carried within a MIME multipart/related message in such a way that the SOAP processing rules for a standard SOAP message are not changed.

An associated W3C document Web services Description Language (WSDL)(<http://www.w3.org/TR/wsdl>) outlines a technique for including bindings to MIME types in a WSDL file.

## Writing the WSDL extensions for SOAP messages with attachments

### Usage scenario

The following WSDL illustrates a simple operation that has one attachment called `attch`:

```
<binding name="MyBinding" type="tns:abc" >
  <soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/http"/>
  <operation name="MyOperation">
    <soap:operation soapAction=""/>
    <input>
      <mime:multipartRelated>
        <mime:part>
          <soap:body parts="part1 part2 ..." use="encoded" namespace="http://mynamespace"
            encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
        </mime:part>
      </mime:multipartRelated>
    </input>
  </operation>
</binding>
```

```

    <mime:part>
      <mime:content part="attch" type="text/html"/>
    </mime:part>
  </mime:multipartRelated>
</input>
</operation>
</binding>

```

**Note:**

- There must be a part (in this example attch) on the message for the operation (in this example MyOperation). There can be other input parts to MyOperation that are not attachments.
- In the binding input there must either be a <soap:body tag or a <mime:multipartRelated tag, but not both.
- For MIME messages, the soap:body is inside a mime:part. There may be one mime:part that contains a soap:body in the binding input and that must not contain a mime:content as well, because a content type of text/xml is assumed for the soap:body.
- There can be multiple attachments in a MIME message, each described by a mime:part.
- Each mime:part (that is not a soap:body) contains a mime:content that describes the attachment itself. The type attribute inside the mime:content is not checked or used by the gateway. It is there to suggest to the application using the gateway what the attachment contains. Multiple mime:contents inside a single mime:part means that the backend service will expect a single attachment with a type specified by one of the mime:contents inside that mime:part.
- The parts="..." attribute inside the soap:body is assumed to contain the names of all the SOAP parts in the message, but not the attachment parts. If there are only attachment parts, then you must specify parts="" (empty string). If you omit the parts attribute altogether, then the gateway assumes ALL parts including the attachments - which means the attachments will appear twice.

In your WSDL you might have defined a schema for the attachment (for instance as a binary[]). Whether or not you have done this, the gateway silently ignores this mapping and treats the attachment as a Data Handler.

Unreferenced attachments need not be mentioned in the WSDL bindings at all.

## Developing Web services gateway extensions

Use this task to develop your own extensions for the gateway

### Before you begin

This information is intended for use by Java programmers.

To extend the functionality of the Web services gateway, you can write your own Java programs. The gateway does not provide any application programming interfaces, but there are system level interfaces that you can use. Specific guidance on how to do this in key areas is given in the following topics:

- Writing a filter for the Web services gateway.
- Using a filter to select a target service and port.
- Capturing Web service invocation information from the Web services gateway.
- Handling exceptions for the Web services gateway.

For additional technical details of the Web services gateway, see the Javadoc ([../javadoc/wsg/index.html](http://../javadoc/wsg/index.html)).

## Writing a filter for the Web services gateway

Use this task to write a filter for the Web services gateway.

### Before you begin

To use this information you should be familiar with using a J2EE session bean development environment such as IBM WebSphere Studio Application Developer.

A Web services gateway filter is essentially a J2EE session bean implementing specific Home and Remote interfaces.

To write a filter using IBM WebSphere Studio Application Developer, complete the following steps. For more detailed information on writing session beans, see the WebSphere Studio Application Developer documentation topic "Developing enterprise beans - overview".

### Steps for this task

1. Open the J2EE perspective.
2. To create a new EJB application project, complete the following steps:
  - a. Select **File -> New -> Enterprise Application Project**.  
The Project Creation wizard opens.
  - b. In the Project Creation wizard, complete the following steps:
    - 1) Select the version of the J2EE specification that you want to use, then click **Next**.
    - 2) Type your project name.
    - 3) Create a new module project for the EJB project only. Clear the other **Application Client Project**, **Web Project** and **Connector Project** check boxes as necessary.
    - 4) Clear the **Web module** check box.
    - 5) Click **Finish**.  
Your new EJB application project is created.
3. To add the extra JAR files that your EJB module needs that are not already in the Enterprise Application Server `/lib` directory, complete the following steps:
  - a. Select **File -> Import**.
  - b. Select the input source **File system**, then click **Next**.
  - c. In the Import window, complete the following steps:
    - 1) Select `WSGW_HOME/client` as the source directory.  
where `WSGW_HOME` is the root directory for your installation of the gateway.
    - 2) Select `wsgwejb.jar`.
    - 3) Select the root directory of your new project as the **destination for imported resources**.
    - 4) Click **Finish**.
  - d. **(Optional)** Repeat the previous **File -> Import** process to add any other extra JAR files that your EJB module needs.
  - e. In the J2EE Hierarchy view, from the pop-up menu for your EJB module, select **Open With -> JAR Dependency Editor**.

- f. In the JAR Dependencies window, select all the JAR files listed.
  - g. Close the JAR Dependencies window, then click **Yes** in the Save Resource window to save your changes.
4. To add extra JAR files to the Java build path for your EJB module, complete the following steps:
    - a. In the J2EE Hierarchy view, select your EJB module's **Properties**.
    - b. In the Properties window, ensure that the following JAR files are included on the **Java Build Path**:
      - *WAS\_HOME/lib/jrom.jar*
      - *WAS\_HOME/lib/qname.jar*
      - *WAS\_HOME/lib/wsd14j.jar*
      - *WAS\_HOME/lib/wsif.jar*
      - *WSGW\_HOME/client/wsgwejb.jar*

where *WAS\_HOME* is the root directory for your installation of WebSphere Application Server Enterprise, and *WSGW\_HOME* is the root directory for your installation of the gateway.

- c. Add any other JAR files or projects that you need for compiling your filter.
  - d. Click **OK**.
5. To create the session bean, complete the following steps:
    - a. Select **File -> New -> Enterprise Bean**.  
The Enterprise Bean Creation wizard opens.
    - b. In the Enterprise Bean Creation wizard, complete the following steps:
      - 1) Select your EJB project, then click **Next**.
      - 2) Ensure that **Session Bean** is selected.
      - 3) Enter a name for the bean.
      - 4) Enter a suitable package name for the bean.
      - 5) Click **Next**.
    - c. In the Enterprise Bean Details window, complete the following steps:
      - 1) Accept the defaults offered for **Session type** (Stateless) and **Transaction type** (Container).
      - 2) Accept the defaults offered for **Bean supertype** (<none>), **Bean class** and **EJB binding name**.
      - 3) Confirm that **Local client view** is not enabled.
      - 4) For the **Remote client view: Remote Home Interface**, click **Class...** then select the `com.ibm.wsgw.beans.FilterHome` interface.
      - 5) For the **Remote client view: Remote Interface**, click **Class...** and select the `com.ibm.wsgw.beans.FilterRemote` interface.
      - 6) Click **Next**.
    - d. In the EJB Java Class window, specify the **Bean superclass** as `com.ibm.wsgw.beans.Filterimpl`, then click **Finish**.

Your new session bean is created.

6. The generated java code for your session bean does not implement the Filter. To update the code, complete the following steps:
  - a. In the J2EE Hierarchy view, expand your session bean to show Java code entries for the Home interface, the Remote interface and for the session bean itself.

- b. In the J2EE Hierarchy view, double-click the entry for the session bean code. In the editor view, the generated code opens for editing.
- c. In the editor view, add the following import statements:

```
import com.ibm.wsgw.*;
import com.ibm.wsgw.beans.*;
import org.apache.wsif.*;
import java.rmi.RemoteException;
```

**Note:** If you selected J2EE version 1.3 or later for your Enterprise Application Project, then omit the statement `import java.rmi.RemoteException` from the set of import statements added here, and see the note below.

- d. Select **File -> Save** to save the file. Ignore any errors at this stage.

**Note:** If you selected J2EE version 1.3 or later for your Enterprise Application Project, then your filter methods must not throw `java.rmi.RemoteException`. So for J2EE version 1.3 or later:

- Omit the statement `import java.rmi.RemoteException` from the set of import statements added in the step above.
- Remove all instances of `throws RemoteException` from your filter methods.
- Code your filter to override the `getContextVersion()` method from `FilterImpl` (in order to eliminate the `RemoteException` from this method). Specifically, implement a `getContextVersion()` method in your filter to return `Filter.CONTEXT_VERSION_WORKAREA` as follows:

```
public int getContextVersion()
{
    return Filter.CONTEXT_VERSION_WORKAREA;
}
```

7. To add the unimplemented methods of the Filter interface to your session bean, complete the following steps:
  - a. Open the Outline view (select **Window -> Show View -> Outline**).
  - b. In the Outline view, from the pop-up menu for your session bean, select **Override Methods**.
  - c. In the Override Methods window, select all the Filter methods to override then click **OK**.

The methods of the Filter interface are added to your session bean.

8. Select **File -> Save** to save the file. Any errors from the previous **File -> Save** are resolved.
9. Develop your filter.

The exact steps that you take to develop your filter depend upon what you want it to do. However to develop any filter, you use the following resources:

- The Filter interface.
- The gateway Javadoc for the Filter interface
- The additional information on the Filter interface that is in Web services gateway - the Filter interface
- The gateway message context. (This contains the context values for each message that comes into the gateway. These are the values that your filter acts upon.)
- The gateway Javadoc for the GatewayContextNames class. (To use the gateway message context values, you import the GatewayContextNames class.)

- The additional information on the gateway message context values that is in Web services gateway - the message context values.
- The gateway WorkArea. (Filters use this to get and set the message context values, as described in Web services gateway - the Filter interface.)
- The WSIF Javadoc for the following WSIF objects:
  - WSIFRequest
  - WSIFResponse
  - WSIFMessage
  - WSIFException

(the methods of the Filter interface use these objects, as described in Web services gateway - the Filter interface.)
- The example code below.

**Note:** You must observe the J2EE programming model, and ensure that any non-gateway services you use are available on all platforms that the filter might be expected to run on. For example, you should not use static variables to store state information because on certain platforms, or in certain configurations (such as a cluster), a filter might be invoked in a different JVM for each request.

### Usage scenario

This example shows you how to access the context and get values in the `filterRequest` method of a filter.

```
import javax.naming.InitialContext;
import javax.naming.NamingException;

import com.ibm.websphere.workarea.UserWorkArea;
import com.ibm.websphere.workarea.WorkAreaException;

import com.ibm.wsgw.GatewayContextNames;

...

try
{
    // Lookup the WorkArea gateway context in JNDI
    InitialContext ctx = new InitialContext();
    UserWorkArea wsgwContext =
        (UserWorkArea)ctx.lookup("services:websphere/WSGW/workarea");

    // Get the currently selected port name
    String Ptype =
        (wsgwContext.get(GatewayContextNames.TARGET_PORT_NAME)).getClass().getName();
    String ThePortname =
        (String) wsgwContext.get(GatewayContextNames.TARGET_PORT_NAME);

    // Get the currently selected target service WSDL location
    String Xtype =
        (wsgwContext.get(GatewayContextNames.TARGET_SERVICE_LOCATION)).getClass().getName();
    TargetServiceLocation WSDLObject =
        (TargetServiceLocation) wsgwContext.get(GatewayContextNames.TARGET_SERVICE_LOCATION);

    String ServiceLocation = WSDLObject.serviceLocation;
    int ServiceLocationType = WSDLObject.serviceLocationType;
    String ServiceName = WSDLObject.serviceName;
    String ServiceNamespace = WSDLObject.serviceNamespace;

}
catch ( NamingException e )
{
```

```

    // Handle any exceptions thrown by the InitialContext here
}
catch ( WorkAreaException e )
{
    // Handle exceptions thrown by UserWorkArea here
}
...

```

### What to do next

After you have developed your filter, you need to generate deployment code and export the enterprise application. To do this using IBM WebSphere Studio Application Developer, complete the following steps:

1. Open the J2EE perspective.
2. In the J2EE Hierarchy view, from the pop-up menu for your EJB module, select **Generate -> Deploy and RMIC code**.
3. In the Generate Deploy and RMIC Code window, select the beans for which you want to generate code, then click **Finish**.
4. To configure the deployment descriptor properties for your bean, complete the following steps:
  - a. In the J2EE Hierarchy view, from the pop-up menu for your bean, select **Open With -> EJB Deployment Descriptor**.
  - b. On the **Beans** tab, set the JNDI name to the Filter class name. This name will be used as the "Home Location" when the filter is deployed to the gateway.
  - c. Close the EJB Deployment Descriptor window, then click **Save** to save the changes.
5. In the J2EE Hierarchy view, from the pop-up menu for your project, select **Export EAR file** to export the enterprise application.

You are now ready to install your filter (as described in the penultimate step of Installing the gateway into a deployment manager cell and Installing the gateway into a stand-alone application server), then deploy your filter.

### **Web services gateway - the Filter interface**

This topic gives more information on using each of the methods of the Filter interface. It supplements the information given in the following Javadoc:

- The gateway Javadoc for the Filter interface
- The WSIF Javadoc for
  - WSIFRequest
  - WSIFResponse
  - WSIFMessage
  - WSIFException

The Filter interface represents an object which is called during service invocation. A bean which implements this interface can be registered to be called just before a request invocation, or just after response receipt for a particular service.

Gateway filters use the gateway WorkArea to get and set the gateway message context values. You get the gateway's WorkArea partition from `services:websphere/WSGW/workarea` in JNDI, then use the UserWorkArea API to get and set data within the message context.

These are the two main methods you use for developing your filter:

- `FilterAction filterRequest(org.apache.wsif.WSIFRequest request, org.apache.wsif.WSIFResponse response).`
- `FilterAction filterResponse(org.apache.wsif.WSIFRequest request, org.apache.wsif.WSIFResponse response).`

Another important method that requires a specific value to be returned is `getContextVersion()`.

If you want your filter to change the `WSIFResponse` and `WSIFRequest` messages, then note that changes to messages are only recognized if the `setUpdatedRequest` and `setUpdatedResponse` methods are called on the returned `FilterAction` object. The `FilterAction` object can also dictate whether processing of the message should continue by calling the `setContinueProcess` method.

### **filterRequest Method**

The `filterRequest` method is called by the gateway Manager just before a request is sent to a target service. The return value from the method can indicate that the request should not be sent.

The request parameter contains the request `WSIFMessage`. This consists of a set of named parts. Each part has a value which is encoded as an instance of an appropriate Java object. Filters can change the values of the Java object instances, but should not add or remove parts, or replace the values of parts with ones of a different type.

The Filter might decide that the request should not proceed. In that case it has three options:

- Throw a `FilterException`. The gateway logs the exception but continues processing filters and the request invocation.
- Throw a `WSGWException`. The gateway logs and rethrows the exception, and processing of filters and the request is stopped. The exception then goes back to the receiving channel, and the channel must determine what to do with the exception (in the case of SOAP-based channels, this results in a Fault message back to the client). This should only be done for unexpected errors in the filter.
- Return a `FilterAction` object with the `continueProcessing` flag set to false. In this case the response message in the `FilterAction` can also be set, and is sent to the originator of the request. No further filters are invoked.

If the request or response is modified, then it must be returned in an instance of the `FilterAction` class. If this is not done, any change to the response is ignored by the Web services gateway.

### **filterResponse Method**

The `filterResponse` method is called by the gateway Manager just after a response has been received from a target service. The response parameter contains the response or fault `WSIFMessage`. This consists of a set of named parts.

Each part has a value which is encoded as an instance of an appropriate Java object. Filters can change the values of the Java object instances, but should not add or remove parts, or replace the values of parts with ones of a different type.



The Filter may decide that the response should not proceed. In that case it has three options:

- Throw a `FilterException`. The gateway logs the exception but continues processing filters and the response invocation.
- Throw a `WSGWException`. The gateway logs and rethrows the exception, and processing of filters and the response is stopped. The exception then goes back to the receiving channel, and the channel must determine what to do with the exception (in the case of SOAP-based channels, this results in a Fault message back to the client). This should only be done for unexpected errors in the filter.
- Return a `FilterAction` object with the `continueProcessing` flag set to false. In this case the fault WSIF Message is set in the response message in the `FilterAction`, and is sent to the originator of the request. No further filters are invoked.

If the response is modified, then it must be returned in an instance of the `FilterAction` class. If this is not done, any change to the response is ignored by the Web services gateway.

If the Filter throws a `FilterException`, it is logged, but the gateway continues to process other filters. If it throws a `WSGWException`, processing of the response is stopped.

See also [Creating and returning a SOAP fault message from a `filterResponse` method](#).

### **`getVersionString` Method**

The `getVersionString` method returns a string form of the version of the filter implementation. This is used by the gateway when logging events relating to the filter so that the exact version of the filter implementation is known.

### **`getContextVersion` Method**

The `getContextVersion` method indicates the approach that this filter uses to access context information. To access the message context information for IBM WebSphere Application Server Version 5, this method must be implemented to return the value: `Filter.CONTEXT_VERSION_WORKAREA`.

### **`init` Method**

The `init` method tells the filter that it has been configured with the Web services gateway.

This method is called by the gateway when it has been asked to add a filter.

### **`destroy` Method**

The `destroy` method tells the filter that it is no longer configured with the Web services gateway. This method is called by the gateway when it has been asked to remove a filter.

## **Creating and returning a SOAP fault message from a `filterResponse` method**

This topic gives an example of the recommended way to create and return a SOAP fault message from a `GatewayFilter.filterResponse` method.

### Usage scenario

```
public FilterAction filterResponse(WSIFRequest wsifRequest, WSIFResponse wsifResponse)
    throws FilterException, WSGWException, RemoteException {
    // Construct the fault message
    WSIFMessage faultMessage = new WSIFDefaultMessage();
    faultMessage.setObjectPart(WSIFConstants.SOAP_FAULT_ACTOR, "mySoapFaultActor");
    faultMessage.setObjectPart(WSIFConstants.SOAP_FAULT_CODE, "mySoapFaultCode");
    faultMessage.setObjectPart(WSIFConstants.SOAP_FAULT_STRING, "mySoapFaultString");
    faultMessage.setObjectPart("stackTrace", "myStackTraceDetails");
    faultMessage.setObjectPart("otherDetails", "myOtherDetails");
    // repeat faultMessage.setObjectPart("aaaa", "bbbb");
    // for each additional detail element
    // Set the fault message into the wsifResponse object
    wsifResponse.setFaultMessage(faultMessage);
    wsifResponse.setIsFault(true);

    // Return the updated response in the filterAction object
    FilterAction filterAction = new FilterAction();
    filterAction.setUpdatedResponse(wsifResponse);
    filterAction.setContinueProcessing(false);
    return filterAction;
}
```

**5.0.2** This example works with both the Apache SOAP and the SOAP/HTTP Channels. Using the SOAP/HTTP channel, the example returns the following SOAP fault response:

```
<?xml version="1.0" encoding="UTF-8"?>
<soapenv:Envelope xmlns:soapenv="http://schemas.xmlsoap.org/soap/envelope/"
  xmlns:xsd="http://www.w3.org/2001/XMLSchema"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance">
  <soapenv:Body>
    <soapenv:Fault>
      <faultcode xmlns:ns1="http://xml.apache.org/axis/">ns1:mySoapFaultCode</faultcode>
      <faultstring>mySoapFaultString</faultstring>
      <faultactor>mySoapFaultActor</faultactor>
      <detail>
        <stackTrace>myStackTraceDetails</stackTrace>
        <otherDetails>myOtherDetails</otherDetails>
      </detail>
    </soapenv:Fault>
  </soapenv:Body>
</soapenv:Envelope>
```

### Web services gateway - the gateway message context values

The gateway message context contains the context information for each incoming message.

You can use the **Context Field Constant** values if you import the class `com.ibm.wsgw.GatewayContextNames`.

For basic information on the fields that are available in the context, see the Javadoc for the `GatewayContextNames` class. Additional information on all of these fields except `AUTH_SUBJECT` and `copyright` is given in the table below.

**Note:** In this version of the gateway you should treat all of the context fields as **Read only**. If your filter attempts to write to a context field, you do not get an error message (because the write does not actually fail) but subsequent system behavior is not readily predictable.

If you want to change the target service location and port name fields, then you should use the Routing interface to get the list of valid target service locations and

to select the target service location. For more information see Using a filter to select a target service and port.

| <b>Context Field Constant</b> | <b>Description</b>  |
|-------------------------------|---|
| AUTH_PASSWORD                 | Read the password from the incoming HTTP request (where available)  |
| AUTH_USER_NAME                | Read the user name from the incoming HTTP request (where available).  |
| GATEWAY_SERVICE_NAME          | Name of the gateway Service for which the request was received.   |
| MESSAGE_ID                    | Set by the channel and is a server-unique ID that can be used to correlate messages, for example in trace. Can be made globally unique by prefixing with host name.                                   |
| RECEIVING_CHANNEL_NAME        | Name of the channel on which the request was received.  |
| REQUEST_SOAP_HEADERS          | Retrieve the SOAP headers for an inbound SOAP request. SOAP Headers are returned as a Vector of Nodes.  |
| RESPONSE_SOAP_HEADERS         | Retrieve the SOAP headers for an outbound SOAP response. SOAP Headers are returned as an Array of Nodes.  |
| RETRY_COUNT                   | Number of retries for the request. NOT CURRENTLY USED.  |
| SOAP_HEADERS                  | Retrieve the SOAP headers for an inbound SOAP request. SOAP Headers are returned as a Vector of Nodes. DEPRECATED (use REQUEST_SOAP_HEADERS).   |
| TARGET_PORT_NAME              | Currently selected port name. Not set until after service invocation, so can only be got by response filters and not by request filters. See also Using a filter to select a target service and port. |
| TARGET_SERVICE_LOCATION       | Gives the location of the currently selected target service's WSDL, Name and Namespace. See also the Javadoc for the TargetServiceLocation class.   |
| TIMEOUT_TIME                  | Time-out value for the response. NOT CURRENTLY USED.  |

## Using a filter to select a target service and port

Use this task to write a filter to select a target service and port.

When a request is received by the gateway, it must determine what the target service really is, and what port to use to access that service.

The gateway represents each exported service as a gateway service. Each gateway service can map to one or more target services, but without filters there is no point in mapping multiple targets as the gateway will always pick the first one. If you want to map multiple targets, you also need to write pluggable filters (configured for each gateway service) that can select the target service from those available.

You write a filter as described in Writing a filter for the Web services gateway. Your filter can get the list of potential target services from the Routing service. It needs to select the target service, then call the Routing service to set the target service (note that doing this clears any prior selection of a target service's port). Your filter might also use the Routing service to select the target port for the service invocation.

The Routing service provides a non-standard interface which is defined in the topic The Routing interface. The implementation of the Routing service is not pluggable.

The Home object for this service must implement the `com.ibm.wsgw.beans.RoutingHome` interface and be located in JNDI at `websphere/WSGW/Routing`.

The sequence of events for a filter to determine and set the target service is as follows:

1. The filter is called with a `WSIFRequest`.
2. The filter obtains the list of potential target services from the Routing service.
3. The filter selects the target service.
4. The filter calls the Routing service to set the target service (note that doing this clears any prior selection of a target service's port).

The Filter can also use the Routing service to select the target port for the service invocation.

Each target service is identified by the target service definition location (which is unique) and target service identity information (which might not be unique). So to select the target service, your filter can either get the table of mappings from target service location to identity information, then choose a target service to use; or it can call `setSelectedTargetServiceIdentity` with the required identity string (relying on the target service identity information being carefully defined). The routing service then selects the first target service it finds (for the current gateway service) with identity information that matches that specified (using `String.equals()`).

**Note:** When you use Routing to set the target service or the target port, the Routing service updates the request context. Because the request context has changed, you then need to set the request object in the `FilterAction` object that you return from the `filterRequest` method (see the filter interface).

### **Web services gateway - the Routing interface**

This topic gives more information on using each of the methods of the Routing interface. It supplements the information given for this interface in the gateway Javadoc.

The Routing interface encapsulates a service which manages routing for requests. Filters can use this service to select the target service and port.

**Note:** The set methods all return a `WSIFRequest` object that contains the updated context information.

To get information on the currently selected target service, use the following Routing interface methods:

#### **getSelectedTargetServiceLocation**

This method returns the currently selected target service location for the request.

#### **getTargetPortName**

This method returns the currently selected target port name for the request.

**Note:** In this version of the gateway, this method always returns blank.

**getTargetServiceDefinition**

This method returns the currently selected target service definition for the request.

To set the target port, use the following Routing interface method:

**setTargetPortName**

This method sets the selected target port name for the request.

To get information on all potential target services, use the following Routing interface method:

**getTargetServices**

This method returns the set of target service names which are mapped by the gateway service on which the request was received.

To set the target service, use either of the following Routing interface methods:

**setSelectedTargetServiceLocation**

This method sets the selected target service location for the request. The selected port name if any is reset by this call.

**setSelectedTargetServiceIdentity**

This method sets the selected target service identity for the request. Target service identity need not be unique, so the first target service found with matching identity information is set. If none is found that matches, the method throws a WSGWException.

## Capturing Web service invocation information from the Web services gateway

Use this task to help you to capture Web service invocation information from the Web services gateway.

The Web services gateway has not implemented a service that stores operational messages, but the gateway does contain an interface (the **MessageWarehouse** interface) to encapsulate such a service. This interface is driven by channels on receipt of requests and before sending responses.

If you have your own system for handling (classifying, storing and retrieving) operational messages, you can potentially use it to log the gateway's operational messages through the MessageWarehouse interface.

The Home object for this service must implement the `com.ibm.wsgw.beans.MessageWarehouseHome` interface and be located in JNDI at `websphere/WSGW/MessageWarehouse`.

### Web services gateway - the MessageWarehouse interface

This topic gives more information on using each of the methods of the MessageWarehouse interface. It supplements the information given for this interface in the gateway Javadoc.

The MessageWarehouse interface encapsulates a service which stores messages for archiving. This interface is used by the channels to log incoming requests for the purposes of non-repudiation.

A default implementation of this interface is not provided by the Web services gateway. If no implementation is present the interface is not used.

### **logRequest**

This method stores a request, along with information about the channel and originator of the request.

It is called by a channel when a request is received, after the user has been authenticated and the message decrypted. The channel may provide information to identify the originator of the request, and to identify the channel itself. The request itself is logged as a `WSIFMessage`.

Additional information regarding receipt of the request (for example any associated digital certificates) can also be logged as `Serializable` objects.

### **logResponse**

This method stores a response, along with information about the channel and destination of the response.

It is called by a channel when a response is about to be sent, before the response is encrypted. The channel may provide information to identify the destination of the response, and to identify the channel itself. The response itself is logged as a `WSIFMessage`.

Additional information regarding sending of the response (for example any associated digital certificates) can also be logged as `Serializable` objects.

### **logException**

This method stores a request, along with information about the channel and originator of the request in the event that an exception is thrown to the channel while the request is being processed. This method allows the exact request and exception information to be logged before the channel decides what actions to take as a result of the exception.

Additional information regarding request and exception (for example any associated digital certificates) can also be logged as `Serializable` objects.

## **Handling exceptions for the Web services gateway**

Use this task to help you to capture information on the gateway's exception handling activities.

During normal processing of a Web service invocation, a fault message might be generated by the target service, and is passed back to the channel to be sent to the originator. As far as the Web services gateway is concerned there is no difference between processing a normal output message and processing a fault message.

But when an exception occurs during processing of a request, the channel needs some way to decide what to do with the exception. What is needed is a service that provides a pluggable handler that can look at the message, exception and other information to decide whether the exception should be thrown back to the originator, or whether a fault message should be constructed.

This service is not provided with the Web services gateway, but the gateway does contain an interface to encapsulate such a service. The `ExceptionHandler` interface allows channels to call an exception handling service, and allows the exceptions to be reported to a third party for analysis.

The Home object for this service must implement the `com.ibm.wsgw.beans.ExceptionHandlerHome` interface and be located in JNDI at `websphere/WSGW/ExceptionHandlerService`.

## Web services gateway - the ExceptionHandler interface

This topic gives more information on using each of the methods of the ExceptionHandler interface. It supplements the information given in the gateway this interface in the gateway Javadoc.

The ExceptionHandler interface encapsulates a service which takes actions when exceptions occur during request/response processing in the gateway. It is intended for use by channels to allow a centralized facility to report and take actions when exceptions occur within the gateway.

### handleException

This method is called by a channel when an exception is caught as a result of processing a message.

The return value indicates what action the channel should take. The actions include:

- Re-throw the original exception.
- Throw a new exception (this is thrown by the handler itself).
- Convert the exception into a fault message.

**Note:** If there is no ExceptionHandler installed, the MessageWarehouse (if any) is always used to log the exception, then the exception is rethrown.

---

## Administering security for the Web services gateway

**5.0.2** The gateway provides facilities for secure communication between the service requester and the gateway, and between the gateway and the target service. To configure these facilities for each deployed Web service, see Enabling Web Services Security (WS-Security) for the gateway.

The gateway provides basic authentication and authorization facilities based upon the broader security features of WebSphere Application Server. See Enabling basic authentication and authorization for the gateway.

The gateway can also invoke Web services that include https:// in their addresses, if you have configured your Java and WebSphere Application Server security properties to allow it. To check your security property settings, see Invoking Web services over HTTPS.

**5.0.2** In many installations, the gateway requires access to the Internet for invoking external Web services. To configure the gateway to work with authenticating proxies, see Enabling proxy authentication for the gateway.

### What to do next

For hints on solving security-related problems, see "Web services gateway troubleshooting tips".

## Enabling Web Services Security (WS-Security) for the gateway

You can configure the gateway for secure transmission of SOAP messages using tokens, keys, signatures and encryption in accordance with the Web Services Security (WS-Security) draft recommendation. For more information on how WS-Security is implemented in WebSphere Application Server Network

Deployment, see (Securing Web services). For more information on the approach taken by the gateway to implementing this emerging standard, see The Web services gateway and WS-Security.

The gateway sits between the service requester (the client) and the target Web service. You configure the gateway to act as the target service from the point of view of the client, and as the client from the point of view of the target service. So you need to get, from the owning parties, the WS-Security configurations for both the client and the Web service. This information is found in the following files on the owners systems:

- Key stores (.ks and .jceks files).
- Certificate stores (.cer files).
- Security settings (ibm-webservicesclient-ext.xmi for the client, and ibm-webservices-ext.xmi for the Web service).
- Binding information - for example the location of a keystore on the file system (ibm-webservicesclient-bnd.xmi for the client, and ibm-webservices-bnd.xmi for the Web service).

**Note:** If the client is hosted on WebSphere Application Server, and the Web service security settings were created using IBM Web services tooling (for example WebSphere Studio Application Developer), then the files that contain the security settings and binding information will have the exact file names (\*.xmi) given above. For clients and Web services from other vendors, these files will have different names.

You need to copy the key store and certificate store files to the gateway file system, and to enter and configure for the gateway the security settings that are contained in the .xmi files. The security settings are entered and configured manually using the gateway administrative user interface. There are tools available (for example WebSphere Studio Application Developer) that can parse the .xmi files for you.

You use the **Gateway -> Security** option to configure the security bindings (the tokens, keys, signatures and encryption methods) that are available to the gateway, as described in Configuring the gateway security bindings.

You then configure the level of security that applies at each stage of the transmission (and note that different levels of security, including no security, can be applied to each stage):

- From the service requester to the gateway.
- From the gateway to the target service.
- From the target service back to the gateway.
- From the gateway back to the service requester.

For information on how to do this, see the following topics:

- Editing the service security configuration - how to configure secure communication for this gateway service between the service requester (the client) and the gateway.
- Editing the target service security configuration - how to configure secure communication between the gateway and the target service.


### **The Web services gateway and WS-Security**

You can configure the gateway for secure transmission of SOAP messages using tokens, keys, signatures and encryption in accordance with the emerging Web



Services Security (WS-Security) specification(<http://www-106.ibm.com/developerworks/webservices/library/ws-secure/>).

In a normal (non gateway) WS-Security scenario, the flow is as shown in the following figure:

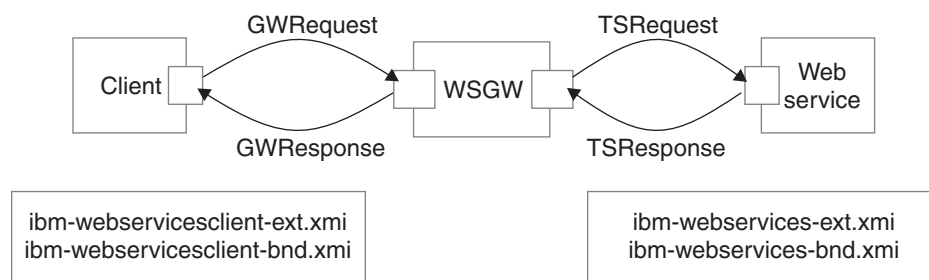


The client generates a request which is then handled by the client Web services engine. It reads the client security configuration and applies the security defined in the `ibm-webservicesclient-ext.xmi` file to the SOAP message. It gets additional binding information from the `ibm-webservicesclient-bnd.xmi` file (for instance, the location of a keystore on the file system).

On receipt of a SOAP message, the Web services engine on the server refers to the `*.xmi` files for the called Web service. In this case, the `ibm-webservices-ext.xmi` file tells the engine what security the incoming message must have (for example, that the body of the message must be signed). If the message does not comply, then it is rejected. The Web services engine verifies any security information, then passes the message on to the Web service being called.

On the response leg from server to client, the process is reversed. The Web service `*.xmi` files tell the Web services engine what security to apply to the response message, and the client `*.xmi` files tell the client engine what security to require in the response message.

When the gateway is introduced, the scenario is more complex. Essentially it can be thought of as two separate request/response invocations. Client to gateway and gateway to target service, as shown in the following figure:



In this scenario, the client and the Web service are unchanged, and still have the same security settings in their `*.xmi` files. However, the gateway is unsecured. Secure SOAP messages cannot travel through the gateway unchanged, and must be processed on receipt. So the gateway needs to act as the target service from the point of view of the client, and as the client from the point of view of the target

service. This scenario means that the security settings for the Web service need to be configured for the view of the service that the gateway presents to the client, and the security settings for the associated gateway target services (remember that there may be multiple target services deployed for a single gateway service) need to be configured with the security settings for the client.

WS-Security settings for the gateway are configured manually using the gateway administrative user interface.

### Configuring the gateway security bindings

You use the **Gateway -> Security** option to configure all the security bindings (the tokens, keys, signatures and encryption methods) that are available to the gateway. This information describes the security bindings that are used to secure the SOAP messages that pass between service requesters (clients) and the gateway, and between target services and the gateway.

You receive this security binding information direct from the service requester or target service provider, in the form of an `ibm-webservicesclient-bnd.xmi` file for the client, and an `ibm-webservices-bnd.xmi` file for the Web service. You extract the information from these `.xmi` files, then manually enter it into the gateway security bindings forms that are described below.

#### Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

#### Gateway

- [Security](#)

The **Configure Gateway Security Bindings** form is displayed.

This form is divided into eight sections, one for each type of security binding. The following comments apply to every section:

- To add a new binding, click **Add new *binding\_type***.
  - To amend an existing binding, click on the name of the binding.
  - To delete an existing binding, click **remove** alongside the name of the binding.
  - If you choose to add or amend a binding, then a binding information sub-form is displayed. Within this form:
    - Help is provided in comments on the sub-form, and in hover-help alongside each field.
    - Fields marked on the sub-form with an asterisk (\*) are required fields.
    - For some bindings, the fields that are required are different depending on whether you are *adding* or *amending* a binding.
    - Many fields are populated by making a selection from a drop-down list.
    - Information entered in one binding information sub-form often appears in a drop-down list in another sub-form.
    - When you have finished entering information for a binding, click **OK**.
3. Add, amend or delete **Signing Information**.

The *signing information* specifies the configuration for digital signature validation and message signing.

If you choose to add or amend a signing information binding, then the **Define Signing Information** form is displayed. In this form you enter a name for the binding (if you are adding a new binding), and the following binding details:

- Signature method
- Digest method
- Canonicalization method
- Signing key name
- Signing Key Locator
- Certificate Path (a choice between trusting *any* certificates, and trusting certificates from a particular *Trust Anchor* and *Certificate Store*).

**Note:** The signing information can also be used for X.509 certificate validation when **Authentication Method** is IDAssertion and **ID Type** is X509Certificate in the gateway service configuration. In such cases, you must fill in the "Certificate Path" fields only.

4. Add, amend or delete **Encryption Information**.

The *encryption information* specifies the configuration for encrypting and decrypting messages.

If you choose to add or amend an encryption information binding, then the **Define Encryption Information** form is displayed. In this form you enter a name for the binding (if you are adding a new binding), and the following binding details:

- Key locator
- Encryption key name
- Key encryption algorithm
- Data encryption algorithm

5. Add, amend or delete **Trust Anchors**.

A *trust anchor* specifies a list of key store configurations that contain root trusted certificates.

If you choose to add or amend a trust anchor, then the **Define Trust Anchor** form is displayed. In this form you enter a name for the trust anchor (if you are adding a new anchor), and the following details:

- Key store type
- Key store path
- Key store password

**Note:** These configurations are used for certificate path validation of the incoming X.509-formatted security tokens. The keystore must be created using the Java Development Kit keytool. The ikeyman tool is not supported.

6. Add, amend or delete **Certificate Stores**.

A *certificate store* specifies a list of untrusted, intermediate certificate files. It is used for certificate path validation of incoming X.509-formatted security tokens.

If you choose to add or amend a certificate store, then the **Define Certificate Store** form is displayed. In this form you enter a name for the certificate store (if you are adding a new store), and the following details:

- Certificate Store Provider

When you amend an existing certificate store, you are given an extra option to add or remove X.509 certificates from the list of certificates that are contained within this store. When you add an X.509 certificate, you specify the full path for the certificate.

7. Add, amend or delete **Key Locators**.

A *key locator* specifies a configuration that is used to retrieve keys for signature and encryption. A key locator class can be customized to retrieve keys from other types of repositories. The default implementation retrieves keys from a keystore.

If you choose to add or amend a key locator, then the **Define Key Locators** form is displayed. In this form you enter a name for the key locator (if you are adding a new key locator), and the following details:

- Classname
- Key store type
- Key store path
- Key store password

When you amend an existing key locator, you are given two extra options:

- Add or remove key entries. For each additional key, you specify:
  - Key name
  - Key alias
  - Key password

**Note:** You do not need to list all the certificate entries as keys; instead, the distinguished name (DN) of the certificate is used as the search key.

- Add or remove additional properties for the configuration. For each additional property, you specify:
  - Property name
  - Property value

8. Add, amend or delete **Trusted ID Evaluators**.

A *trusted ID evaluator* determines whether the identity (ID)-asserting authority is trusted.

If you choose to add or amend a trusted ID evaluator, then the **Define Trusted ID Evaluator** form is displayed. In this form you enter a name for the trusted ID evaluator (if you are adding a new evaluator), and the following details:

- Class name

When you amend an existing trusted ID evaluator, you are given an extra option to add or remove additional properties for the configuration. For each additional property, you specify:

- Property name
- Property value

9. Add, amend or delete **Login Mappings**.

A *login mapping* specifies a configuration for validating security tokens within incoming messages.

If you choose to add or amend a login mapping, then the **Define Login Mapping** form is displayed. In this form you enter a name for the login mapping (if you are adding a new mapping), and the following details:

- Authentication method

- JAAS Configuration name
- Callback handler factory class name
- Token type local name
- Token type URI

When you amend an existing login mapping, you are given extra options to add or remove additional properties for the *configuration*, and additional properties for the *callback handler factory*. For each additional property, you specify:

- Property name
- Property value

#### 10. Add, amend or delete **Login Bindings**.

A *login binding* specifies a configuration for generating security tokens within outgoing messages.

If you choose to add or amend a login binding, then the **Define Login Binding** form is displayed. In this form you enter a name for the login binding (if you are adding a new binding), and the following details:

- Authentication method
- JAAS Configuration name
- Callback handler factory class name
- Token type local name
- Token type URI

When you amend an existing login binding, you are given extra options to add or remove additional properties for the *configuration*, and additional properties for the *callback handler*. For each additional property, you specify:

- Property name
- Property value

## **Editing the service security configuration**

### Before you begin

Before you can select the security settings that are applied for an individual Web service, you must configure the gateway security bindings.

For each Web service, you can select the security settings that are applied between the service requester (the client) and the gateway. These settings are specified for each stage of the transmission:

- From the service requester to the gateway (the client request).
- From the gateway back to the service requester (the gateway response).

You receive this security settings information from the service requester and from the target service provider in the following form:

- An `ibm-webservicesclient-bnd.xml` for the client, and `ibm-webservices-bnd.xml` for the Web service, from which you process the security bindings information as described in *Configuring the gateway security bindings*.
- An `ibm-webservicesclient-ext.xml` for the client, and `ibm-webservices-ext.xml` for the Web service, which contain the information on the levels of security (integrity, confidentiality and identification) that are required when this Web

service exchanges messages with a service requester. These are therefore also the settings that the gateway needs to apply when it makes the equivalent gateway service available to a service requester.

To set the security settings that are applied between the service requester (the client) and the gateway, complete the following steps:

Steps for this task

1. List the gateway-deployed Web services
2. Click the name of a Web service in the list.  
A form is displayed through which you can view and modify the current deployment details for this Web service, and add or remove multiple target services.
3. In the **Service Security** section, select the **Edit service security configuration** option.

The service security configuration form is displayed. This form is divided into the following sections:

- **Gateway Security Properties** (the Actor URI)
- **Client Request Security Properties** (integrity, confidentiality and identification settings)
- **Gateway Response Security Properties** (the response Actor URI, and integrity and confidentiality settings)
- **Security bindings** (request bindings and response bindings).

The following comments apply to every section:

- Help is provided in comments on the form, and in hover-help alongside each field.
  - There are no required fields.
  - Many fields are populated by making a selection from a drop-down list.
4. In the **Gateway Security Properties** section, set the Actor URI.  
**Note:** If you specify an Actor URI, then only SOAP security headers with this Actor URI will be processed.
  5. In the **Client Request Security Properties** section, set the following security levels:
    - a. Set the **Integrity** level.  
Set the parts of the incoming SOAP message that must be signed (the Body, the Timestamp and the Security Token).
    - b. Set the **Confidentiality** level.  
Set the parts of the incoming SOAP message that must be encrypted (the Body and the Username Token).
    - c. Set the **Identification** level.  
Set the identification methods that will be accepted (Basic Authentication , Digital Signature and ID Assertion).
  6. In the **Gateway Response Security Properties** section, set the following security levels:
    - a. Set the Response Actor URI.  
**Note:** If you specify a Response Actor URI, then the SOAP security header in the response message will have this Actor URI.
    - b. Set the **Integrity** level.

Set the parts of the response SOAP message that must be signed (the Body and the Timestamp).

- c. Set the **Confidentiality** level.

Set whether or not the Body of the response SOAP message must be encrypted.

7. In the **Security bindings** section, set the Request bindings and the Response bindings that are to be used.

**Note:** You choose these bindings (Signing Information, Encryption Information, Trusted ID Evaluator and Login Mappings) from pull-down lists. The available items in these lists are those that you have previously entered as described in Configuring the gateway security bindings.

8. When you have finished editing the service security configuration, click **Apply changes**.

## Editing the target service security configuration

### Before you begin

Before you can select the security settings that are applied for target Web service, you must configure the gateway security bindings.

For each target Web service, you can select the security settings that are applied between the target Web service and the gateway. These settings are specified for each stage of the transmission:

- From the gateway to the target service (the target service request).
- From the target service back to the gateway (the target service response).

You receive this security settings information from the service requester and from the target service provider - usually in the following form:

- An `ibm-webservicesclient-bnd.xmi` for the client, and `ibm-webservices-bnd.xmi` for the Web service, from which you process the security bindings information as described in Configuring the gateway security bindings.
- An `ibm-webservicesclient-ext.xmi` for the client, and `ibm-webservices-ext.xmi` for the Web service, which contains the information on the levels of security (integrity, confidentiality and identification) that are required when this Web service exchanges messages with a service requester. These are therefore also the settings that the gateway needs to apply when it calls the target service on behalf of the service requester.

To set the security settings that are applied between the target service and the gateway, complete the following steps:

### Steps for this task

1. List the gateway-deployed Web services
2. Click the name of a Web service in the list.

A form is displayed through which you can view and modify the current deployment details for this Web service, and add or remove multiple target services.

3. In the **Target Services** section, click the name of a target Web service in the list.

A form is displayed, containing the same fields that you filled in when you added the target service, and also the following additional fields:

- **Started** (a check box).
- **Enable target service security** (a check box).

- **Edit target service security configuration.**
4. Select the **Edit target service security configuration** option.  
The target service security configuration form is displayed. This form is divided into the following sections:
    - **Target Service Security Properties** (the Actor URI)
    - **Target Service Request Security Properties** (integrity, confidentiality and identification settings)
    - **Target Service Response Security Properties** (the response Actor URI, and integrity and confidentiality settings)
    - **Security bindings** (request bindings and response bindings).

The following comments apply to every section:

- Help is provided in comments on the form, and in hover-help alongside each field.
  - There are no required fields.
  - Many fields are populated by making a selection from a drop-down list.
5. In the **Target Service Security Properties** section, set the Actor URI.  
**Note:** If you specify an Actor URI, then only SOAP security headers with this Actor URI will be processed.
  6. In the **Target Service Request Security Properties** section, set the following security levels:
    - a. Set the Target Actor URI.  
**Note:** If you specify a Target Actor URI, then the SOAP security header in the request message will have this Actor URI.
    - b. Set the **Integrity** level.  
Set the parts of the outgoing SOAP message that must be signed (the Body, the Timestamp and the Security Token).
    - c. Set the **Confidentiality** level.  
Set the parts of the outgoing SOAP message that must be encrypted (the Body and the Username Token).
    - d. Set the **Identification** level.  
Set the identification methods that will be accepted (Basic Authentication , Digital Signature and ID Assertion).
  7. In the **Target Service Response Security Properties** section, set the following security levels:
    - a. Set the **Integrity** level.  
Set the parts of the response SOAP message that must be signed (the Body and the Timestamp).
    - b. Set the **Confidentiality** level.  
Set whether or not the Body of the response SOAP message must be encrypted.
  8. In the **Security bindings** section, set the Request bindings and the Response bindings that are to be used.  
**Note:** You choose these bindings (Signing Information, Encryption Information, Trusted ID Evaluator and Login Mappings) from pull-down lists. The available items in these lists are those that you have previously entered as described in Configuring the gateway security bindings.
  9. When you have finished editing the target service security configuration, click **Apply changes**.



## Enabling basic authentication and authorization for the gateway

**5.0.2** In addition to the security options described in Enabling Web services Security (WS-Security) through the gateway, you can also use the broader security features of WebSphere Application Server to enable basic authentication and authorization.

### 5.0+

#### 5.0.1+

The Web services gateway provides a basic authentication and authorization mechanism based upon the broader security features of WebSphere Application Server.

Basic authentication can be applied at two levels, as described in the following topics:

1. Enabling gateway-level authentication.
2. Enabling Web service operation-level authorization.

For **gateway-level authentication**, you set up a role and realm for the gateway on WebSphere Application Server's Web server and servlet container, and define the userid and password that is used by the gateway to access the role and realm. You also modify the gateway's channel applications so that they only give access to the gateway to service requesters that supply the correct userid and password for that role and realm.

**Note:** This means that gateway-level authentication must be enabled before you install any channels.

For **operation-level authorization**, you apply security to individual methods in a Web service. To do this, you create an enterprise bean with methods matching the Web service operations. These EJB methods perform no operation and are just entities for applying security. Existing WebSphere Application Server authentication mechanisms can be applied to the enterprise bean. Before any Web service operation is invoked, a call is made to the EJB method. If authorization is granted, the Web service is invoked. Your target Web service is protected by wrapping it in an EAR file, and applying role-based authorization to the EAR file. This process is explained in general terms in Operation-level security - role-based authorization.

**Note:**

- If you want to enable operation-level authorization, you must first enable gateway-level authentication.
- If you want to change the default gateway-level authentication settings, you must do so before you install any channels.
- After gateway-level authentication has been enabled, filters have access to the requester's authentication information.

The Web services gateway can also invoke Web services that include https:// in their addresses, if the Java and WebSphere security properties have been configured to allow it. To check your security property settings, see the following topic:

- Invoking Web services over HTTPS

## What to do next

For hints on solving security-related problems, see "Web services gateway troubleshooting tips".

### **Enabling gateway-level authentication**

A number of default gateway-level authentication settings are included in the gateway. There is a default role of `AuthenticatedUsers` which includes the special group `'AllAuthenticatedUsers'`. When security is enabled, you must supply a user ID and password to use the gateway administrative interface or invoke a gateway service.

This task covers the three main areas in which you might want to make changes:

- Changing the default gateway-level authentication settings.
- Enabling gateway-level authentication.
- Assigning users and groups to roles.

#### **Note:**

- If you want to change the default gateway-level authentication settings, you must do so before you install any channels. When you run the script that installs the gateway itself (*either* into a deployment manager cell *or* into a stand-alone application server) you also install the following channels:
  - Apache SOAP channel 1.
  - **5.0.2** SOAP/HTTP channel 1.

So if you change the default gateway-level authentication settings after you install the gateway, you then need to re-run the gateway install.

- You can enable gateway-level authentication, and assign users and groups to roles, at any time.
- After gateway-level authentication has been enabled, filters have access to the requester's authentication information.

#### Steps for this task

1. To change the default gateway-level authentication settings, use the WebSphere Application Server (Tool (AAT)) to complete the following steps:
  - a. Set up a role and realm for the gateway on WebSphere Application Server's Web server and servlet container.
  - b. Define the user ID and password that is used by the gateway to access the role and realm.
  - c. Modify the gateway's channel applications so that they only give access to the gateway to service requesters that supply the correct user ID and password for that role and realm.
2. To enable gateway-level authentication, complete the following steps:
  - a. Start the WebSphere Application Server administrative server.
  - b. Start the administrative console.
  - c. In the navigation pane, select **Security -> Global Security**.
  - d. In the main pane, on the **Configuration** tab, enable the "Enabled" check box.
  - e. Save the settings.
  - f. Stop then restart the application server.
  - g. Close the administrative console.

3. You can use the AAT or the administrative console to assign users and groups to roles. To map users to roles using the administrative console, complete the following steps:
  - a. Start the WebSphere Application Server administrative server.
  - b. Start the administrative console.
  - c. In the navigation pane, select **Application -> Enterprise Applications -> wsgw**.  
In the main pane, an option to map security roles to users and groups appears in the Additional Properties table.
  - d. Modify the security roles and save the settings.
  - e. Repeat the previous two steps for each enterprise application that you want to modify.
  - f. Stop then restart the application server.
  - g. Close the administrative console.

For more information see Assigning users and groups to roles.

**Note:** The current jacl install scripts do not let you assign users to roles as part of installing the gateway into a deployment manager cell or into a stand-alone application server.

#### What to do next

You might now want to enable operation-level authorization, or install the gateway.

### **Enabling operation-level authorization**

Use this task to apply security to individual methods in a Web Service.

#### Before you begin

Before you begin this task you must first enable gateway-level authentication.

You can only apply operation-level authorization to a Web service that has already been deployed to the gateway with the check box 'Authorization Policy - Control access to this service' enabled.

This task involves making changes to the file `/lib/wsgwauth.ear`. To protect the installation version of this file, you should make a backup copy of it before you change it.

For operation-level authorization you create an enterprise bean with methods matching the Web service operations. These EJB methods perform no operation and are just entities for applying security. Existing WebSphere Application Server authentication mechanisms can be applied to the enterprise bean. Before any Web service operation is invoked, a call is made to the EJB method. If authorization is granted, the Web service is invoked.

Your target Web service is protected by wrapping it in an EAR file, and applying role-based authorization to the EAR file. This process is explained in general terms in Operation-level security - role-based authorization.

The EAR file that now contains your Web service is then imported into `wsgwauth.ear` (which contains all of the gateway's protected Web services) and

wsgwauth.ear is modified to set the roles and assign them to methods. Finally, this modified wsgwauth.ear file is deployed in WebSphere Application Server and users are assigned to the previously defined roles.

To enable Web service operation-level authorization, complete the following steps:

#### Steps for this task

1. To create *your\_webservice.ear*, complete the following steps:

- a. Open a command prompt.
- b. Go to directory /WSGW/scripts/auth
- c. Enter the command `WSGWAAuthGen location your_webservice` where
  - *location* is the URL for the gateway (this must include the root context)
  - *your\_webservice* is the name of the service as deployed in the gateway (this is case-sensitive)

For example `WSGWAAuthGen http://host:port/wsgw AddressBook` where *host* and *port* are the host name and port number for the application server on which the gateway is installed.

**Note:** The Web service name and operation name can contain characters (such as "-", ".", "&") that are disallowed in an EJB class name and method name. So these are translated during the generation process of *your\_webservice.ear*. A message appears informing you of the name change.

*your\_webservice.ear* is created in directory /WSGW/scripts. There is also a temporary directory /WSGW/scripts/ejb, which you can delete.

2. To finish assigning roles and protecting methods, use the WebSphere Application Server (Application Assembly Tool (AAT)) to complete the following steps:

- a. Start the AAT.
- b. From the File menu select **File > Open**, and browse to select file /lib/wsgwauth.ear.
- c. To import *your\_webservice.ear* into wsgwauth.ear, complete the following steps:
  - In the navigation pane, open the pop-up menu for **EJB Modules** and select **Import**
  - Browse to select file /WSGW/scripts/*your\_webservice.ear*. The Select modules to import window opens.
  - In the Select modules to import window, select *your\_webservice* and click **OK**.
  - The Confirm values window opens. Click **OK**.
  - In the navigation pane, expand **EJB Modules** to confirm that *your\_webservice.ear* has been imported.
- d. In the navigation pane, expand **EJB Modules > your\_webservice.ear** and select **Security Roles**.
- e. For every security role that you want to create, repeat the following steps:
  - From the pop-up menu for **Security Roles**, select **New**.
  - Type the name and description of the new security role, and click **OK**.

- f. In the navigation pane, expand **EJB Modules** > *your\_webservice.ear* and select **Method Permissions**.
  - g. For every defined role that you want to assign to a Web service method, repeat the following steps:
    - From the pop-up menu for **Method Permissions**, select **New**. The New Method Permission window opens.
    - Type the name of the new method permission, and click **ADD** for Methods. The Add Methods window opens.
    - In the Add Methods window, expand the tree for remote methods and select the method to be protected. Click **OK**. The Add Methods window closes.
    - In the New Method Permission window, click **ADD** for Roles. Select a previously defined role from the list then click **OK**.
  - h. To ensure that the authorization enterprise bean can reference the newly-imported enterprise bean, complete the following steps:
    - In the navigation pane, expand **WSGW Authorization group** > **Session Beans** > **Authorization** and select **EJB References**.
    - From the pop-up menu for **EJB References**, select **New**. The New EJB Reference window opens.
    - In the New EJB Reference window, on the **General** tab, type a name for the reference then use the 'Link' combination box to select the newly-imported EJB (all the other fields on this tab are populated automatically).
    - In the New EJB Reference window, on the **Bindings** tab, type the JNDI name as it appears in the bindings tab of the service enterprise bean (this should be in the form *websphere/WSGW/Security/your\_webservice*).
    - Click **OK**. The New EJB Reference window closes.
  - i. From the AAT File menu, select **File** > **Generate Code For Deployment**.
  - j. Make a note of the name of the modified ear file, then click **Generate Now**.
  - k. From the AAT File menu, select **File** > **Save** to save the modified copy of *wsgwauth.ear*.
  - l. Close the AAT.
3. To install the modified copy of *Deployed\_wsgwauth.ear*, complete the following steps:
    - a. Start the WebSphere Application Server Administrative Console.
    - b. In the navigation pane, select **Applications** > **Install an Application**.
    - c. Use **Install New Application** to install *Deployed\_wsgwauth.ear*. Select the users or groups to be assigned to the roles when prompted.

### Operation-level security - role-based authorization

During construction of an EAR file, roles can be defined and applied to methods. At deployment of the EAR file, individual users or groups can be assigned to roles. So you can use this feature of EAR files to add role-based security to your Web service.

For example: You have a Web service that controls access to important information, and you want to give read-only access to some users, and write access to others. So when you build the EAR file you define two roles READ and WRITE, then you apply the READ role to the *getData* method and the WRITE role to the *writeData* method. When you deploy the EAR file in WebSphere Application Server, you assign 'All Authenticated Users' to the READ role and individual users to the

WRITE role. When a user tries to access `WebService.getData`, their user name and password is checked by the operating system or by Lightweight Third Party Authentication (LTPA).

## Invoking Web services over HTTPS

The Web services gateway can invoke Web services that include `https://` in their addresses, if the Java and WebSphere security properties have been configured to allow it. This means that one gateway can send a SOAP/HTTPS message direct to another gateway, rather than having to export services and have clients invoke them using HTTPS.

To enable your gateway to send and receive SOAP/HTTPS messages, confirm that your Java and WebSphere security properties are configured as described in the following steps:

### Steps for this task

1. Check that there is a copy of file `ibmjsse.jar` in directory `WAS_HOME/java/jre/lib/ext` (where `WAS_HOME` is the root directory for your installation of IBM WebSphere Application Server).
2. Edit the security properties file `WAS_HOME/java/jre/lib/security/java.security` so that it includes entries for both the Sun security provider and the IBM security provider. For example:  

```
security.provider.1=sun.security.provider.Sun
security.provider.2=com.ibm.jsse.IBMJSSEProvider
```

**Note:** The order is significant. The Sun security provider must come before the IBM provider.

3. Use the WebSphere Application Server Administrative Console to set up the following equivalent system properties:

```
// truststore location
System.setProperty("javax.net.ssl.trustStore",
    "your_truststore_root_directory/TestSSL/key.jks");
// set truststore password
System.setProperty("javax.net.ssl.trustStorePassword",
    "your_truststore_password");
//use ibm reference implementation
System.setProperty("java.protocol.handler.pkgs",
    "com.ibm.net.ssl.internal.www.protocol");
```

## Enabling proxy authentication for the gateway

The gateway requires access to the Internet for invoking Web services and for retrieval of WSDL files. Many enterprise installations use a proxy server in support of Internet routing, and many proxy servers require authentication before they grant access to the Internet. This requirement is supported in HTTP messaging by a "Proxy-Authorization" message header that contains encoded username and password credentials.

For messages passing through the gateway, you can enable and disable proxy authentication, and specify whether the authentication credentials are supplied by the service requester or by the gateway. If you specify *requester-supplied* credentials, the credentials in the HTTP message that the gateway receives are re-instantiated by the gateway in the equivalent message that it sends on to the proxy. If you specify *gateway-supplied* credentials, the gateway ignores any credentials in the incoming HTTP message and supplies its own credentials in the equivalent message that it sends on to the proxy.

**Note:** In certain circumstances, the gateway also creates and sends its own messages (for example for WSDL retrieval). In these cases, the gateway always supplies its own credentials to the authenticating proxy. So even if you enable proxy authentication and specify *requester-supplied* credentials, you must still supply credentials for the gateway.

To enable proxy authentication for the gateway, complete the following steps:

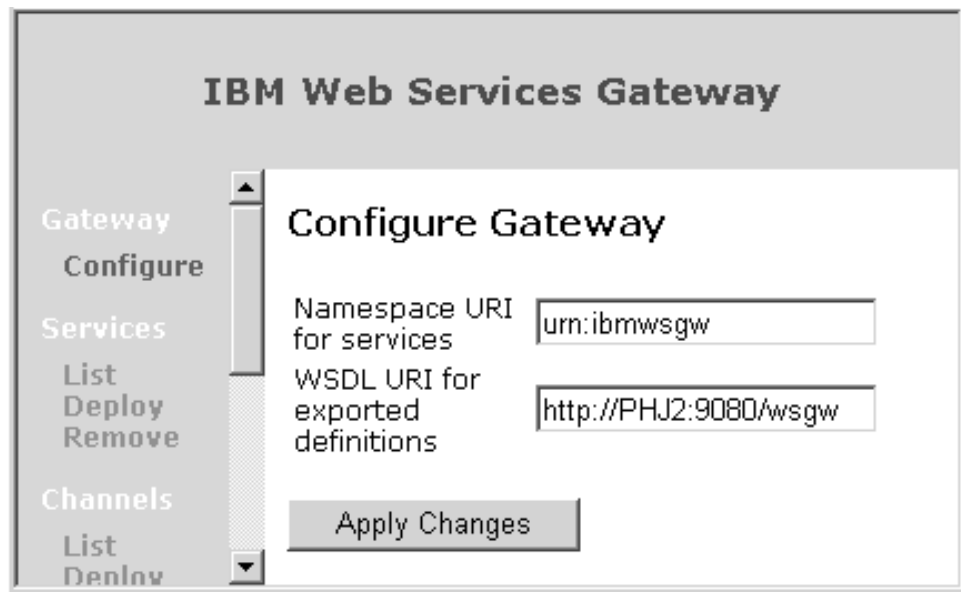
Steps for this task

1. Display the Web services gateway administrative user interface.
2. In the navigation pane, click the following link:

**Gateway**

- Configure

The gateway configuration form is displayed:



The screenshot shows the IBM Web Services Gateway administrative interface. The main title is "IBM Web Services Gateway". On the left is a navigation pane with sections: "Gateway" (containing "Configure"), "Services" (containing "List", "Deploy", "Remove"), and "Channels" (containing "List", "Deploy"). The "Configure Gateway" form is displayed in the main area. It has two text input fields: "Namespace URI for services" with the value "urn:ibmwsgw" and "WSDL URI for exported definitions" with the value "http://PHJ2:9080/wsgw". Below these fields is an "Apply Changes" button.

3. Enable the **Enable proxy authentication** check box.
4. In the **Proxy user** field, type the proxy username for the gateway itself.  
**Note:** If you enable proxy authentication then this field is compulsory, even if you also specify *requester-supplied* credentials as described below.
5. In the **Proxy password** field, type the associated proxy password for the gateway itself.  
**Note:** If you enable proxy authentication then this field is compulsory, even if you also specify *requester-supplied* credentials as described in the next step.
6. To set the **Use Gateway proxy credentials for invoking WebServices** check box, complete one of the following two steps:
  - a. If you want to use *requester-supplied* credentials, then clear the **Use Gateway proxy credentials** check box.  
With this setting, each incoming message to the gateway from a service requester is expected to contain a valid "Proxy-Authorization" HTTP message header. This header is re-instantiated by the gateway in the equivalent message that it sends on to the proxy.

**Note:** For gateway-initiated messaging, such as WSDL retrieval, the gateway supplies its own credentials in the HTTP messages that it sends to the proxy.

- b. If you want to use *gateway-supplied* credentials, then enable the **Use Gateway proxy credentials** check box.

With this setting, a trust association is established between the gateway and the authenticating proxy. The gateway supplies its own credentials in all messages that it sends to the proxy, and no username or password is required from service requesters for invoking Web services.

7. Click **Apply Changes**.
8. You also need to provide the application server in which your gateway is running with machine details for the authenticating proxy and for any internal machines that do not require authentication. You do this by setting system properties in the WebSphere Application Server Java Virtual Machine as follows:
  - a. Start the WebSphere Application Server administrative server.
  - b. Start the administrative console.
  - c. In the navigation pane, select **Application Servers -> your\_server\_name -> Process Definition -> Java Virtual Machine -> Custom Properties**.
  - d. Set the following properties:
    - **http.proxySet** - Set this to true to tell the application server that it is required to work with an authenticating proxy.
    - **http.proxyHost** - Set this to the machine name of the authenticating proxy.
    - **http.proxyPort** - Set this to the port through which the authenticating proxy is accessed. For example 8080
    - **http.nonProxyHosts** - List the internal machines for which authentication is not required for routing through the proxy. Separate each machine name in the list with a vertical bar "|".

**Note:** This list must include the machine on which the gateway is installed.
  - e. Save the settings.
  - f. Stop then restart the application server.
  - g. Close the administrative console.

#### What to do next

**Note:** You also use the gateway configuration form to set the namespace URI and WSDL URI for the Web services gateway.

---

## Web services gateway troubleshooting tips

This topic provides hints to help you resolve problems you experience when using the Web services gateway.

For information on resolving WebSphere-level problems, see *Diagnosing and fixing problems*.

To identify and resolve gateway-related problems, you can use the standard WebSphere Application Server trace and logging facilities. To enable trace for the gateway, set the application server's trace string to `com.ibm.wsgw.*=all=enabled`. If you encounter a problem that you think might be related to the gateway, you can



check for error messages in the WebSphere Application Server administrative console, and in the application server's stdout.log file. You can also enable the application server debug trace to provide a detailed exception dump.

**5.0.2** If you report a gateway-related problem to IBM Service and Support, please include in your report the gateway release and build information that is listed under the **About** option of the Web services gateway administrative user interface.

The gateway's user interface uses cascading style sheets to lay out its pages, and javascript to monitor progress and advise you as you fill in each on-screen form. So your Web browser must support javascript and cascading style sheets, and it must be configured so that javascript and style sheets are enabled. How you do this depends on which browser you use. For example for Netscape, you select **Edit** -> **Preferences**, click **Advanced** in the Category pane, then confirm that the **Enable Javascript** and **Enable style sheets** check boxes are selected.

A list of the gateway runtime system messages, with details of what each message means, is given in Web services gateway messages.

Here is a checklist of common problems:

**You are about to apply a WebSphere Application Server fix pack or upgrade, but you do not want to lose your current gateway configuration.** **5.0.2**

When you apply an upgrade or fix pack to WebSphere Application Server, the configured gateway is replaced with an upgraded but empty gateway. To preserve your gateway configuration, use the **save** application to save the configuration before you apply the upgrade, then use the **restore** application to restore the configuration after you apply the upgrade. For more information, see Backing up and restoring a gateway configuration.

**You have managed to deploy your Web service in the Web services gateway but you are getting a class cast exception when you invoke the operation which takes an integer parameter.**

Check that your client is using the version of soap.jar that is supplied in the WebSphere Application Server's /AppServer/lib/app directory. If you enable trace, you may see in the trace for the request <SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/" xmlns:xsi="http://www.w3.org/1999/XMLSchema-instance" xmlns:xsd="http://www.w3.org/1999/XMLSchema">

The gateway expects the 2001 version of the XML schema. Older versions of soap.jar (including 2.2) generate 1999 schema. If you have the soap.jar that is supplied with WebSphere Application Server in the client's class path, you should see 2001 schema in the request, which should then work.

**The persistent state of the Web services gateway has become out of sync with the channel applications.**

This can happen if you remove and reinstall the Apache SOAP applications. If you need to do this, then either ensure that all corresponding channels configured with the Web services gateway are removed, or remove and reinstall the Web services gateway at the same time.

**Note:** The Web services gateway application (wsgw.ear) must be installed before channel and filter applications. If the gateway application needs to

be reinstalled, all channels and filters must be uninstalled first, then reinstalled after the gateway application.

**You are getting SOAP fault messages, but cannot determine the precise problem from the fault message.**

If you receive a SOAP fault message with a faultstring which is just the value of one of the parameters of the invocation, that means that that parameter's value was invalid. For example if you have a service which expects an int parameter and you send it a message containing the value "1.1", then the fault message you receive simply contains 1.1 as the fault string:

```
<faultcode>SOAP-ENV:Client</faultcode>
<faultstring>1.1</faultstring>
```

**Note:** This is Apache SOAP behavior, and not something that the gateway can do anything about.

If you receive a SOAP fault message containing an element that is not present in the WSDL for the target service, then the error message thrown can be difficult to identify. There are two possible scenarios:

- The WSDL is deployed to use **Generic Classes**. In this case the returned SOAP message contains an `IllegalArgumentException` exception. For example:

```
[Attributes={}] [faultCode=SOAP-ENV:Server]
[faultString=com.ibm.wsgw.WSGWException:
WSGW0043E: Exception while executing
operation createEntry service ExchangeService.
Exception: org.apache.wsif.WSIFException: SOAPException:
SOAP-ENV:ClientNo mapping found for 'com.ibm.jrom.JROMValue'
using encoding style 'http://schemas.xmlsoap.org/soap/encoding/';
nested exception is:
[SOAPException: faultCode=SOAP-ENV:Client; msg=No mapping found for
'com.ibm.jrom.JROMValue' using encoding style
'http://schemas.xmlsoap.org/soap/encoding/';
targetException=java.lang.IllegalArgumentException:
No mapping found for 'com.ibm.jrom.JROMValue' using encoding style
'http://schemas.xmlsoap.org/soap/encoding/'.]]
[faultActorURI=/wsgwsoap1/soaprpcrouter]
...
```

- The WSDL is deployed to use **Deployed Classes**. In this case an empty message is returned. For example:

```
[Attributes={}] [faultCode=null] [faultString=null] [faultActorURI=null]
[DetailEntries=] [FaultEntries=]
```

**Note:** This is Apache SOAP behavior, and not something that the gateway can do anything about.

**You are enabling operation-level authorization, but when you install wsgwauth.ear, an error message appears in the WebSphere Application Server administrative console detailing a JNDI problem.**

Check that you entered, in the authorization session bean's 'EJB References', the correct JNDI name of the imported Web service enterprise bean. Note that this is case sensitive.

**You are trying to have a Web services gateway send an SOAP/HTTPS message to another Web services gateway, and you are receiving a Malformed URLErrorException error.**

The Web services gateway can invoke Web services that include https:// in their addresses, if the Java and WebSphere security properties have been

configured to allow it. To check your security property settings, see the topic *Invoking Web services over HTTPS*

**You deselect 'Authorization Policy - Control access to this service' from the deployment details for a Web service, and you find the service no longer works.**

A number of tasks are required to disable security. Clearing the check box 'Authorization Policy - Control access to this service' will still leave WebSphere Application Server security in place, so basic authentication might still be required.

To disable security fully, use the WebSphere Application Server administrative console's Security Center to disable Global Security.

**You experience problems with handling Document style SOAP messages.**

**5.0 +**

**5.0.1 +**

The SOAP message format must be RPC style. The gateway does not support Document style SOAP messages.

**5.0.2**

If you are using the Apache SOAP Channel, then the SOAP message format must be RPC style. To handle Document style SOAP messages, use the SOAP/HTTP channel (which supports both RPC style and Document style SOAP messages).

**You experience problems with handling SOAP messages with attachments.**

**5.0 +**

**5.0.1 +**

The gateway does not support SOAP messages with attachments.

**5.0.2**

To handle SOAP messages with attachments, use the SOAP/HTTP channel.

**You experience problems with a Web service that uses complex data types. The same service works fine, when not using the gateway.**

You need to do one of two things to support Web services that use complex types in the Web services gateway:

- Set the message part representation for the service to **Generic classes**.
- Set the message part representation for the service to **Deployed Java classes**, and make the original classes available to the application server (either by updating the class path or by putting the JAR file somewhere like *WAS\_HOME/lib/app*). If you do not have the original classes, you can use *wsdl2java* to generate java beans that can contain values of the complex type for the service, then compile the beans into a JAR file and make it available to the application server.

For more information on the factors to consider when choosing between these options, see *Data type representation - choosing between Generic classes and Deployed Java classes*.

**Note:**

- If your Web service has non-bean parameters (that is, it requires a custom serializer/deserializer) then it is not supported by the current release of the gateway.
- Not all complex types that are expressible in XML schema are supported by the current release of the gateway. For more information see Web services gateway - Supported types.
- See also Complex data types - mapping namespaces to packages.

**5.0 +****5.0.1 +**

If your Web service contains complex types, and it is deployed to Axis 1.0, and you used the Axis tools to generate the WSDL file, then you need to remove the following line in the WSDL schema definition:

```
<import namespace="any_namespace_name">
```

## Web services gateway messages

WebSphere system messages are logged from a variety of sources, including application server components and applications. Messages logged by application server components and associated IBM products start with a unique message identifier that indicates the component or application that issued the message. For more information about the message identifier format, see the topic Message Format.

The rest of this topic contains a list of the Web services gateway runtime system messages, with details of what each message means.

**WSGW0001E: Channel name {0} from gateway configuration differs from that in JNDI: {1}**

**Explanation:** The name specified for the channel does not match the name of the channel as defined within the EAR file.

**User Response:** Ensure that the channel name is specified correctly

**WSGW0002E: Error storing endpoint address. Exception: {0}**

**Explanation:** An unexpected exception occurred when storing the endpoint address for a channel.

**User Response:** Contact IBM Support

**WSGW0003E: Error retrieving endpoint address. Exception: {0}**

**Explanation:** An unexpected exception occurred when retrieving the endpoint address for a channel.

**User Response:** Contact IBM Support

**WSGW0004E: Not used**

**Explanation:**

**User Response:**

**WSGW0005E: Error retrieving channel name. Exception: {0}**

**Explanation:** An unexpected exception occurred when retrieving the channel name from JNDI.

**User Response:** Contact IBM Support

**WSGW0006E: Error deploying service to {1}. Exception: {0}**

**Explanation:** An unexpected error occurred deploying the service to the given component.

**User Response:** This error may be caused by a previous failure. Try redeploying the service using a different gateway service name. If that fails, reinstalling the channel and gateway applications may remove the problem.

**WSGW0007E: Error getting endpoint URL from channel {0}. Exception: {1}**

**Explanation:** An unexpected error occurred generating the endpoint URL for the given channel.

**User Response:** Contact IBM Support

**WSGW0008E: Could not determine default port name for target service {0}**

**Explanation:** There are no ports in the WSDL defined for the target service that are supported by currently available WSIF providers or there is an error in the WSDL file associated with the port definition or a namespace it uses.

**User Response:** Either ensure that a WSIF provider is correctly configured for the port in the WSDL, or ensure that the WSDL contains correctly specified port information.

**WSGW0009E: Failed to deploy service. Exception: {0}**

**Explanation:** An unexpected error occurred trying to deploy the service.

**User Response:** Contact IBM Support

**WSGW0010E: The namespaceURI attribute cannot be changed when there are active services**

**Explanation:** The namespaceURI is used to generate WSDL for gateway services. If this global setting is changed then current WSDL becomes invalid.

**User Response:** Either remove all channels or all gateway services from the gateway configuration and retry the change.

**WSGW0011E: Not used**

**Explanation:**

**User Response:**

**WSGW0012E: Not used**

**Explanation:**

**User Response:**

**WSGW0013E: Could not locate home {0}. Exception: {1}**

**Explanation:** The specified home location could not be found in JNDI.

**User Response:** Ensure that the home location is specified correctly, and that it appears in JNDI.

**WSGW0014E: Not used**

**Explanation:**

**User Response:**

**WSGW0015E: Could not create instance of class {0}. Exception: {1}**

**Explanation:** The gateway failed to create an instance of the specified Java class.

**User Response:** Ensure that the Java class has a public constructor with no parameters.

**WSGW0016E: Could not locate class {0}. Exception: {1}**

**Explanation:** The gateway failed to locate the specified Java class.

**User Response:** Ensure that the Java class is visible to the gateway application's classloader.

**WSGW0017E: Not used**

**Explanation:**

**User Response:**

**WSGW0018E: Not used**

**Explanation:**

**User Response:**

**WSGW0019E: Failed to clone definition. Exception: {1}**

**Explanation:** An unexpected error occurred cloning a WSDL definition.

**User Response:** Contact IBM Support

**WSGW0020E: Error while loading mapped type class {0}. Exception: {1}**

**Explanation:** An error occurred while trying to load the given Java class which represents a type in the deployed WSDL for a target service.

**User Response:** Ensure that the Java class is visible to the gateway application's classloader.

**WSGW0021E: Expected WSDL definition to contain a <wsdl:type> element with a schema from one of the '{0}', '{1}', or '{2}' namespaces**

**Explanation:** Schema types in WSDL definitions must be declared using one of the specified XML Schema namespaces.

**User Response:** Update the WSDL definition to use the appropriate namespace.

**WSGW0022E: Unexpected Schema->Java problem when parsing WSDL file.**

**Exception: {0}**

**Explanation:** An unexpected exception occurred when parsing a WSDL file. This may be due to unsupported elements in the WSDL.

**User Response:** Contact IBM Support

**WSGW0023E: Unexpected Schema->JROM problem when parsing WSDL file.**

**Exception: {0}**

**Explanation:** An unexpected exception occurred when parsing a WSDL file. This may be due to unsupported elements in the WSDL.

**User Response:** Contact IBM Support

**WSGW0024E: Channel {0} cannot be removed because it is being used by a deployed service**

**Explanation:** Channels can only be removed when they are not in use by gateway services.

**User Response:** Remove the channel from gateway services to which it is deployed before removing the channel.

**WSGW0025E: Target service identity cannot be specified as null**

**Explanation:** A target service can only be selected using a non-null valid for the identity.

**User Response:** Modify the calling code to ensure that the target service identity value is never null.

**WSGW0026E: Invalid gateway service name {0}. The name must be a valid XML Schema NCNAME.**

**Explanation:** The name specified for the gateway service does not conform to the required definition.

**User Response:** Correct the gateway service name so that it is a valid XML Schema NCNAME.

**WSGW0027E: Port {0} does not exist for target service {1}**

**Explanation:** The requested port does not exist for the target service.

**User Response:** Ensure that a valid port is requested, or update the target service WSDL to contain a port of the requested name.

**WSGW0028E: No binding for port {0} for target service {1}**

**Explanation:** The requested port for the target service does not have a binding defined in the WSDL definition of the service.

**User Response:** Ensure that the target service WSDL has a binding for the requested port, or use a different port name.

**WSGW0029E: No portType for binding {0} for port {1} for target service {2}**

**Explanation:** The requested port for the target service does not have a portType defined in the WSDL definition of the service.

**User Response:** Ensure that the target service WSDL has a portType for the requested port, or use a different port name.

**WSGW0030E: Not used**

**Explanation:**

**User Response:**

**WSGW0031E: Channel name {0} already exists**

**Explanation:** The name specified for the channel is the same as that of a channel that is currently deployed.

**User Response:** Choose a different name for the channel, or remove the existing channel of the given name.

**WSGW0032E: Channel name {0} not found**

**Explanation:** No channel is currently deployed with the given name.

**User Response:** Use the name of a channel that is currently deployed.

**WSGW0033E: Filter {0} cannot be removed because it is being used by a deployed service**

**Explanation:** Filters can only be removed when they are not in use by gateway services.

**User Response:** Remove the filter from gateway services to which it is deployed before removing the filter.

**WSGW0034W: Invocation of filter {0} failed. Exception: {1}**

**Explanation:** An unexpected exception was thrown during processing of the given filter.

**User Response:** Contact IBM Support

**WSGW0035E: Filter context version {0} not supported**

**Explanation:** The context version that the filter requires is not supported by this version of the gateway.

**User Response:** Ensure that the filter is requesting the correct context version. It may be necessary to upgrade the gateway to support the filter.

**WSGW0036E: Target service identity information {0} not matched for gateway service {1}**

**Explanation:** A target service was requested by identity, but the identity information does not match any currently deployed target service.

**User Response:** Ensure that the identity information is correct, and that there is a target service deployed to the given gateway service with the right identity information.

**WSGW0037E: Filter name {0} already exists**

**Explanation:** The name specified for the filter is the same as that of a filter that is currently deployed.

**User Response:** Choose a different name for the filter, or remove the existing filter of the given name.

**WSGW0038E: Filter name {0} not found**

**Explanation:** No filter is currently deployed with the given name.

**User Response:** Use the name of a filter that is currently deployed.

**WSGW0039E: Error loading state from {0}. Exception {1}**

**Explanation:** An unexpected exception occurred loading the state of the gateway from the given location.

**User Response:** Ensure that the given location is visible to the gateway application.

**WSGW0040E: Failed to convert definition to string. Exception: {0}**

**Explanation:** An unexpected exception occurred converting a WSDL definition into a string in order to display it or return it to an application.

**User Response:** Contact IBM Support

**WSGW0041E: Failed to save state. Exception {0}**

**Explanation:** An unexpected exception occurred when saving the state of the gateway.

**User Response:** Contact IBM Support

**WSGW0042W: No target services available to get service definition**

**Explanation:** A request was made for the WSDL definition for the gateway service, however no target services have been defined for the gateway service, so it is not possible to generate a WSDL definition.

**User Response:** Deploy one or more target services to the gateway service.

**WSGW0043E: Exception while executing operation {0} service {1}. Exception: {2}**

**Explanation:** An unexpected exception occurred when passing a request on to a target web service.

**User Response:** Ensure that the gateway service and target service are correctly deployed (using the correct message part representation). Ensure that the target service is available and responds correctly to direct requests (i.e. not through the gateway).

**WSGW0044E: Filter position {0} invalid**

**Explanation:** The specified position for addition or removal of the filter was not valid.

**User Response:** Ensure a valid value is specified. The value should be -1, 0 or a positive integer.



**WSGW0045E: Filter not found in list**

**Explanation:** An attempt was made to remove a filter from a gateway service specifying -1 as the index, but the filter is not in the list at all.

**User Response:** Ensure that the correct filter is specified.

**WSGW0046E: Channel {0} already defined for gateway service {1}**

**Explanation:** The given channel has already been defined for the gateway service.

**User Response:** Ensure that the correct channel name is specified.

**WSGW0047E: Channel {0} not defined for gateway service {1}**

**Explanation:** The channel cannot be removed from the gateway service as it is not currently defined for the gateway service.

**User Response:** Ensure that the correct channel name is specified.

**WSGW0048E: UDDI reference {0} already defined for gateway service {1}**

**Explanation:** The given UDDI reference has already been defined for the gateway service.

**User Response:** Ensure that the correct UDDI reference name is specified.

**WSGW0049E: UDDI reference {0} not defined for gateway service {1}**

**Explanation:** The UDDI reference cannot be removed from the gateway service as it is not currently defined for the gateway service.

**User Response:** Ensure that the correct UDDI reference name is specified.

**WSGW0050E: Target service with location {0} already defined for gateway service {1}**

**Explanation:** The given target service location has already been defined for the gateway service.

**User Response:** Ensure that the correct target service location is specified.

**WSGW0051E: Target service with location {0} not defined for gateway service {1}**

**Explanation:** The target service location cannot be removed from the gateway service as it is not currently defined for the gateway service.

**User Response:** Ensure that the correct target service location is specified.

**WSGW0052E: Target service with location {0} was not found for gateway service {1}**

**Explanation:** The target service WSDL definition could not be obtained from the given location.

**User Response:** Ensure that the correct target service location is specified.

**WSGW0053E: gateway service {0} cannot be removed as active entities and force not specified**

**Explanation:** A gateway service with one or more target services, channels, filters or UDDI references cannot be removed.

**User Response:** Remove the target services, channels, filters and UDDI references from the gateway service.

**WSGW0054E: An exported definition for gateway service {0} is not available as there are no defined channels for the service**

**Explanation:** A request was made for the WSDL definition for the gateway service, however no channels have been defined for the gateway service, so it is not possible to generate a WSDL definition.

**User Response:** Deploy one or more channels to the gateway service.

**WSGW0055E: Not used**

**Explanation:**

**User Response:**

**WSGW0056E: No default target service available for {0}**

**Explanation:** The default target service location cannot be obtained for the gateway service as no target services are defined.

**User Response:** Ensure that one or more target services are defined for the gateway service.

**WSGW0057E: No receiving channel name in context**

**Explanation:** A request has reached the gateway that does not contain the receiving channel name in the context.

**User Response:** Contact the supplier of the channel application.

**WSGW0058E: Channel {0} not defined for gateway service {1}**

**Explanation:** A request has reached the gateway for the given service through a channel which is not defined for that service. The request is rejected.

**User Response:** If the channel should be valid for the service, add the channel, otherwise check that the client of the request is making a valid request. This exception may be thrown when a client is making a malicious attack.

**WSGW0059E: gateway service {0} does not exist**

**Explanation:** A request was made for a gateway service that does not exist.

**User Response:** Ensure that the correct gateway service name is specified.

**WSGW0060E: gateway service {0} already exists**

**Explanation:** An attempt was made to create a new gateway service using a name that is used by an existing gateway service.

**User Response:** Use a different name for the gateway service.

**WSGW0061E: Could not find Service in UDDI registry {0} with parameters {1}, {2}, {3}**

**Explanation:** The given parameters for UDDI lookup did not yield a match.

**User Response:** Ensure that the parameters are correct. Also ensure that the UDDI reference parameters are correct and correspond to those used to publish the service to UDDI.

**WSGW0062E: Target service WSDL contains no <service> elements**

**Explanation:** The target service WSDL could be loaded but does not contain a <service> element. This is necessary to be able to invoke the target service.

**User Response:** Ensure that the target service WSDL contains one or more <service> element.

**WSGW0063E: Target service WSDL contains more than one service, and either target service name or namespace not specified**

**Explanation:** When adding a target service to a gateway service, you must specify both the service name and namespace values if there is more than one <service> element in the target service WSDL.

**User Response:** Specify the target service name and namespace as well as the location.

**WSGW0064E: Target service name {0} does not match service name in WSDL: {1}**

**Explanation:** A target service name was specified that is not the same as any target service name in the WSDL at the given location.

**User Response:** Ensure that a valid target service name is specified.

**WSGW0065E: Target service namespace {0} does not match service namespace in WSDL: {1}**

**Explanation:** A target service namespace was specified that is not the same as any target service namespace in the WSDL at the given location.

**User Response:** Ensure that a valid target service namespace is specified.

**WSGW0066E: Target service name {0} or namespace {1} not found in WSDL definition**

**Explanation:** A target service name and namespace were both specified, but do not match any target service name and namespace combination in the WSDL at the given location.

**User Response:** Ensure that a valid target service name and namespace combination is specified.

**WSGW0067E: UDDI reference {0} cannot be removed because it is being used by a deployed service**

**Explanation:** UDDI references can only be removed when they are not in use by gateway services.

**User Response:** Remove the UDDI reference from gateway services to which it is deployed before removing the UDDI reference.

**WSGW0068E: UDDI reference {0} already exists**

**Explanation:** The name specified for the UDDI reference is the same as that of a UDDI reference that is currently deployed.

**User Response:** Choose a different name for the UDDI reference, or remove the existing UDDI reference of the given name.

**WSGW0069E: UDDI reference {0} not found**

**Explanation:** No UDDI reference is currently deployed with the given name.

**User Response:** Use the name of a UDDI reference that is currently deployed.

**WSGW0070E: Invalid target service location type {0}**

**Explanation:** The location type for the target service is not a valid value.

**User Response:** Ensure that a correct value is specified for the target service location type.

**WSGW0071E: Failed to load URL definition from {0}**

**Explanation:** The URL location specified was incorrect, or the WSDL it refers to cannot be loaded.

**User Response:** Ensure that the URL location is correct, and refers to a valid WSDL document.

**WSGW0072E: Failed to load UDDI definition from {0}**

**Explanation:** The UDDI location specified was incorrect, or the WSDL it refers to cannot be loaded.

**User Response:** Ensure that the UDDI location is correct, and refers to a valid WSDL document.

**WSGW0073W: Not used**

**Explanation:**

**User Response:**

**WSGW0074E: Not used**

**Explanation:**

**User Response:**

**WSGW0075E: Failed to set gateway end point address. Exception {0}**

**Explanation:** An unexpected exception occurred when automatically setting the gateway's end point address.

**User Response:** Contact IBM Support

**WSGW0076E: Unable to access the gateway configuration bean. Exception {0}**

**Explanation:** An unexpected exception occurred looking up the gateway's configuration bean in JNDI.

**User Response:** Restart the application server.

**WSGW0077E: Failed to remove gateway configuration session. Exception {0}**

**Explanation:** An unexpected exception occurred removing the session bean while access the gateway's configuration bean.

**User Response:** Contact IBM Support

**WSGW0078E: Unable to access the gateway EndPoint bean. Exception {0}**

**Explanation:** An unexpected exception occurred looking up the gateway's endpoint bean in JNDI.

**User Response:** Restart the application server.

**WSGW0079E: Failed to remove endpoint session. Exception {0}**

**Explanation:** An unexpected exception occurred removing the session bean while access the gateway's endpoint bean.

**User Response:** Contact IBM Support

**WSGW0080E: Performance monitoring error. Exception {0}**

**Explanation:** An unexpected exception occurred when recording performance monitoring information.

**User Response:** Contact IBM Support

**WSGW0081E: Unexpected error in method {0}. Exception {1}**

**Explanation:** An unexpected exception occurred in the given method.

**User Response:** Contact IBM Support

**WSGW0082E: Unable to determine WAS security setting**

**Explanation:** The WAS security setting could not be determined. It will be assumed that security is enabled.

**User Response:** No action required.

**WSGW0083W: Failed to authorize request for operation {0} on service {1}.  
Exception {2}**

**Explanation:** Authorization of the given request failed. The request has been rejected.

**User Response:** Ensure that the required authorization bean has been generated for the given service, and that the correct authorization policy is defined.

**WSGW0084W: Invocation of filter {0} version {1} failed. Exception {2}**

**Explanation:** An exception was thrown during processing of the given filter. Processing of the request continues.

**User Response:** Investigate the reason for the exception being thrown by the filter. Refer to the documentation for the filter on how to resolve the problem.

**WSGW0085E: Failed to publish service {0} to UDDI registry {1}. Exception: {2}**

**Explanation:** An unexpected exception occurred when publishing the given service to a UDDI registry.

**User Response:** Ensure that the properties of the gateway service and UDDI reference are specified correctly.

**WSGW0086E: Failed to unpublish service {0} from UDDI registry {1}. Exception: {2}**

**Explanation:** An unexpected exception occurred when unpublishing the given service from a UDDI registry.

**User Response:** Ensure that the properties of the gateway service and UDDI reference are specified correctly.

**WSGW0087I: Published service {0} to UDDI registry {1}**

**Explanation:** The service was successfully published to the UDDI registry.

**User Response:** None

**WSGW0088I: Unpublished service {0} from UDDI registry {1}**

**Explanation:** The service was successfully unpublished from the UDDI registry.

**User Response:** None

**WSGW0089I: No MessageWarehouse registered. Requests will not be logged**

**Explanation:** A MessageWarehouse implementation was not found at the expected location in JNDI, so none is being used.

**User Response:** If a MessageWarehouse has been implemented, ensure that it is bound to JNDI at the correct location.

**WSGW0090I: No ExceptionHandler registered. Exceptions will not be handled**

**Explanation:** An ExceptionHandler implementation was not found at the expected location in JNDI, so none is being used.

**User Response:** If an ExceptionHandler has been implemented, ensure that it is bound to JNDI at the correct location.

**WSGW0091I: Usage: java -jar GenAuth -DWAS\_HOME=<was.install.directory> <HostName> <ServiceName>**

where <was.install.directory> is the location of the WebSphere installation directory and <HostName> is the url pointed to by the installation of the gateway and <ServiceName> is the name of the deployed gateway service. (Please note the ServiceName is case sensitive).

For example

```
java -jar GenAuth.jar -DWAS_HOME=  
c:\\websphere\\AppServer http://host.machine.name.com/wsgw ServiceName
```

Successful execution will generate a file named <ServiceName>.ear

**Explanation:** Usage statement. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0092I: Retrieving Service :**

**Explanation:** Progress message indicating that the service definition is being retrieved. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0093I: Retrieving Port Type :**

**Explanation:** Progress message indicating that the port type information is being retrieved. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0094I: Retrieving Methods :**

**Explanation:** Progress message indicating that method information is being retrieved. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0095I: Making Directory :**

**Explanation:** Progress message indicating that a directory is being created. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0096I: Using Directory :**

**Explanation:** Progress message indicating that a directory is being used. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0097I: About to compile....**

**Explanation:** Progress message indicating that a compilation is about to start. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0098I: Command Status :**

**Explanation:** General command status message. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0099I: About to create jar....**

**Explanation:** Progress message indicating that a JAR file is about to be created. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0100I: About to create ear....**

**Explanation:** Progress message indicating that an EAR file is about to be created. This message is used by the WSGWAuthGen command line utility.

**User Response:** No action required.

**WSGW0101E: Error retrieving port from service {1}**

**Explanation:** An error occurred retrieving the port from the service in the WSDL. This message is used by the WSGWAuthGen command line utility.

**User Response:** Ensure that the service name is specified correctly and is deployed to the gateway with at least one target service and one channel.

**WSGW0102E: Error retrieving service {0}**

**Explanation:** An error occurred retrieving the service. This message is used by the WSGWAuthGen command line utility.

**User Response:** Ensure that the service name is specified correctly and is deployed to the gateway with at least one target service and one channel.

**WSGW0103E: Exception while retrieving service definition from URL: {0}/ServiceDefinition?name={1}. Exception: {2}**

**Explanation:** An unexpected exception occurred retrieving WSDL from the given location. This message is used by the WSGWAuthGen command line utility.

**User Response:** Ensure that the service name is specified correctly and is deployed to the gateway with at least one target service and one channel.

**WSGW0104E: Error retrieving methods from service {0}**

**Explanation:** An unexpected exception occurred retrieving the methods that correspond to operations on the service.

**User Response:** Contact IBM Support

**WSGW0105E: Error retrieving WAS\_HOME environment variable**

**Explanation:** The value of the WAS\_HOME environment variable could not be retrieved.

**User Response:** Ensure that the WAS\_HOME variable is set correctly in the environment under which the command is being executed.

**WSGW0106E: Error compiling files**

**Explanation:** An unexpected error occurred compiling the generated Java files.

**User Response:** Contact IBM Support

**WSGW0107E: Error executing JAR command**

**Explanation:** An unexpected error occurred generating a JAR file.

**User Response:** Contact IBM Support

**WSGW0110E: A client attempted to load imported URL {0} for gateway service {1}. This URL is not imported by the definition for that service.**

**Explanation:** An attempt was made to use the gateway's import mapping servlet to load information from a URL that does not correspond to one that is referenced by the WSDL definition for that service.

**User Response:** Ensure that the client is making a valid request. This may be a malicious attempt to obtain information that the client does not have access to.

**WSGW0111W: Unsupported elements within the WSDL definition for target service {0} were ignored. The functionality of this service may be compromised.**

**Explanation:** In order to be able to use the given WSDL definition within the gateway, certain elements of the definition were ignored.

**User Response:** Refer to the service provider's documentation to determine whether this will affect the use of the service.

**WSGW0120E: Exception while removing ConversationPart {0} from Correlation Service. Exception {1}**

**Explanation:** An unexpected exception occurred when using the Correlation Service.

**User Response:** Contact IBM Support

**WSGW0121E: Exception while accessing ConversationPart {0} from Correlation Service. Exception {1}**

**Explanation:** An unexpected exception occurred when using the Correlation Service.

**User Response:** Contact IBM Support

**WSGW0122E: Exception while storing Serializable {0} at Correlation Service.  
Exception {1}**

**Explanation:** An unexpected exception occurred when using the Correlation Service.

**User Response:** Contact IBM Support

**WSGW0123E: Exception while storing Serializable {0} with id {1} at Correlation Service. Exception {1}**

**Explanation:** An unexpected exception occurred when using the Correlation Service.

**User Response:** Contact IBM Support

---

## Web services gateway: Resources for learning

Use the following links to find supplementary information about getting started with the Web services gateway. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

- WebSphere Version 5 Web services Handbook. This Redbook illustrates with suitable demonstration applications how Web services can be implemented using the IBM product portfolio, especially WebSphere Application Server Version 5 and WebSphere Studio Application Developer Version 5. It includes chapters on the gateway and on Web service security.
- Samples Central. The gateway samples, and documentation on how to use them, are available through the Samples Central page of the IBM WebSphere Developer Domain Web site.
- InfoCenter for WebSphere Application Server Edge components. This InfoCenter contains a library of PDF online books covering all aspects of the WebSphere Application Server Edge components. gateway clustering builds upon the load balancing capabilities of these components.
- The Web Services Security (WS-Security) specification. A major area of gateway security is based upon this emerging standard.
- The IBM Web services gateway: Technical Overview. A different version of the gateway is available as a component of a product called IBM WebSphere Business Connection. This brief technical summary from WebSphere Business Connection applies equally well to the version of the gateway in WebSphere Application Server.

For supplementary information about Web services in general, see “Web services: Resources for learning” on page 479.

The gateway builds on the (Web Services Invocation Framework) (WSIF), which allows the gateway to pass on Web service invocations to any WSDL-defined Web service. For supplementary information about WSIF, see “WSIF: Resources for learning” on page 549.



---

## Chapter 12. Class loading

Class loaders affect the packaging of applications and the run-time behavior of packaged applications deployed on application servers.

### Steps for this task

1. Read about class loaders. The article "Class loading: Resources for learning" points to additional sources.
2. If necessary, migrate class-loader Module Visibility Mode settings for Version 4.0.x applications to Version 5.0 application or WAR class-loader policies.
3. When assembling an enterprise application resource (EAR) file that has EJB modules, set the classpath for the class loader to use during packaging.
4. **(Optional)** If an application module uses a resource, create a resource provider that specifies the directory name of the resource drivers. Do not specify the resource JAR file names. All JAR files in the specified directory will be added into the classpath of the WebSphere Application Server extensions class loader.
5. Configure class loaders of an application server for the run-time environment.
  - a. Click **Servers > Application Servers > *server\_name*** and, on the settings page for an application server, set the application class-loader policy and application class-loader mode.

The application class-loader policy controls the isolation of applications running in the system. When set to SINGLE, applications are not isolated; a single application class loader is used to contain all EJB modules, dependency JAR files, and shared libraries in the system. When set to MULTIPLE, applications are isolated from each other; each application receives its own class loader to load that application's EJB modules, dependency JAR files, and shared libraries.

The application class-loader mode specifies the class-loader mode when the application class-loader policy is SINGLE. PARENT\_FIRST causes the class loader to first delegate the loading of classes to its parent class loader before attempting to load the class from its local classpath. PARENT\_LAST causes the class loader to first attempt to load classes from its local classpath before delegating the class loading to its parent. This allows an application class loader to override and provide its own version of a class that exists in the parent class loader.

- b. On the settings page for an application server, click **Classloader**. On the Classloader page, click **New**.
- c. On the settings page for a class loader, specify the class-loader mode. PARENT\_FIRST causes the class loader to delegate the loading of classes to its parent class loader before attempting to load the class from its local classpath. PARENT\_LAST causes the class loader to attempt to load classes from its local classpath before delegating the class loading to its parent. Then, click **OK**.
- d. **(Optional)** On the settings page for a class loader, click **Libraries**. From the Library Ref page, click **Add**. On the settings page for a library reference, specify variables for the library reference as needed and click **OK**. Repeat the previous step until you define a library reference instance for each library file that your application needs. To define a library reference, you must first define one or more shared libraries.

6. When configuring an installed enterprise application for deployment in the run-time environment, set the class-loader mode and the WAR class-loader policy.
7. When configuring an installed Web module for deployment in the run-time environment, set the class-loader mode.

---

## Class loaders

Class loaders are part of the Java virtual machine (JVM) code and are responsible for finding and loading class files. Class loaders affect the packaging of applications and the run-time behavior of packaged applications deployed on application servers.

The run-time environment of WebSphere Application Server uses the following class loaders to find and load new classes for an application in the following order:

1. The **bootstrap, extensions, and CLASSPATH class loaders** created by the JVM.

The bootstrap class loader uses the boot classpath (typically classes in `jre/lib`) to find and load classes. The extensions class loader uses the system property `java.ext.dirs` (typically `jre/lib/ext`) to find and load classes. The CLASSPATH class loader uses the CLASSPATH environment variable to find and load classes.

The CLASSPATH class loader contains the J2EE APIs of the WebSphere Application Server product (inside `j2ee.jar`). Because the J2EE APIs are in this class loader, you can add libraries that depend on J2EE APIs to the classpath system property to extend a server's classpath. However, a preferred method of extending a server's classpath is to add a shared library.

2. A **WebSphere-specific extensions class loader**.

The WebSphere extensions class loader loads the WebSphere run-time and J2EE classes that are required at run time. The extensions class loader uses a `ws.ext.dirs` system property to determine the path used to load classes. Each directory in the `ws.ext.dirs` classpath and every JAR file or ZIP file in these directories is added to the classpath used by this class loader.

The WebSphere extensions class loader also loads resource provider classes into a server if an application module installed on the server refers to a resource that is associated with the provider and if the provider specifies the directory name of the resource drivers.

3. One or more **application module class loaders** that load elements of enterprise applications running in the server.

The application elements can be Web modules, EJB modules, resource adapters, and dependency JAR files. Application class loaders follow J2EE class-loading rules to load classes and JAR files from an enterprise application. The WebSphere run time enables you to associate a shared library classpath with an application.

Each class loader is a child of the class loader above it. That is, the application module class loaders are children of the WebSphere-specific extensions class loader, which is a child of the CLASSPATH Java class loader. Whenever a class needs to be loaded, the class loader usually delegates the request to its parent class loader. If none of the parent class loaders can find the class, the original class loader attempts to load the class. Requests can only go to a parent class loader; they cannot go to a child class loader. If the WebSphere class loader is requested to find a class in a J2EE module, it cannot go to the application module class loader to find that class and a `ClassNotFoundException` occurs. Once a class is loaded by a

class loader, any new classes that it tries to load reuse the same class loader or go up the precedence list until the class is found.

### **Class-loader isolation policies**

The number and function of the application module class loaders depends on the class-loader policies specified in the server configuration. Class loaders provide multiple options for isolating applications and modules to enable different application packaging schemes to run on an application server.

Two class-loader policies control the isolation of applications and modules:

#### **Application class-loader policy**

Application class loaders consist of EJB modules, dependency JAR files, resource adapters, and shared libraries. Depending on the application class-loader policy, an application class loader can be shared by multiple applications (SINGLE) or unique for each application (MULTIPLE). The application class-loader policy controls the isolation of applications running in the system. When set to SINGLE, applications are not isolated. When set to MULTIPLE, applications are isolated from each other.

#### **WAR class-loader policy**

By default, Web module class loaders load the contents of the WEB-INF/classes and WEB-INF/lib directories. The application class loader is the parent of the Web module class loader. You can change the default behavior by changing the application's WAR class-loader policy.

The WAR class-loader policy controls the isolation of Web modules. If this policy is set to APPLICATION, then the Web module contents also are loaded by the application class loader (in addition to the EJB files, RAR files, dependency JAR files, and shared libraries). If the policy is set to MODULE, then each web module receives its own class loader whose parent is the application class loader.

**Note:** WebSphere server class loaders never load application client modules.

For each application server in the system, you can set the application class-loader policy to SINGLE or MULTIPLE. When the application class-loader policy is set to SINGLE, then a single application class loader loads all EJB modules, dependency JAR files, and shared libraries in the system. When the application class-loader policy is set to MULTIPLE, then each application receives its own class loader used for loading that application's EJB modules, dependency JAR files, and shared libraries.

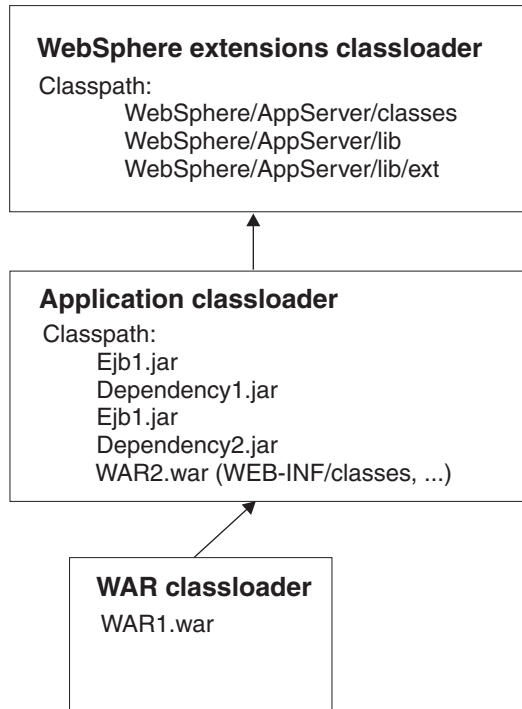
This application class loader can load each application's Web modules if that WAR module's class-loader policy is also set to APPLICATION. If the WAR module's class-loader policy is set to APPLICATION, then the application's loader loads the WAR module's classes. If the WAR class-loader policy is set to MODULE, then each WAR module receives its own class loader.

The following example shows that when the application class-loader policy is set to SINGLE, a single application class loader loads all EJB modules, dependency JAR files, and shared libraries of all applications on the server. The single application class loader can also load Web modules if an application has its WAR class-loader policy set to APPLICATION. Applications having a WAR class-loader policy set to MODULE use a separate class loader for Web modules.

Application class-loader policy: SINGLE

Application 1  
Module: EJB1.jar  
Module: WAR1.war  
MANIFEST Class-Path: Dependency1.jar  
WAR Classloader Policy = MODULE

Application 2  
Module: EJB2.jar  
MANIFEST Class-Path: Dependency2.jar  
Module: WAR2.war  
WAR Classloader Policy = APPLICATION




The following example shows that when the application class-loader policy of an application server is set to **MULTIPLE**, each application on the server has its own class loader. An application class loader also loads its Web modules if the application's WAR class-loader policy is set to **APPLICATION**. If the policy is set to **MODULE**, then a Web module uses its own class loader.

Application class-loader policy: MULTIPLE

Application 1  
Module: EJB1.jar  
Module: WAR1.war  
MANIFEST Class-Path: Dependency1.jar  
WAR Classloader Policy = MODULE

Application 2  
Module: EJB2.jar  
MANIFEST Class-Path: Dependency2.jar  
Module: WAR2.war  
WAR Classloader Policy = APPLICATION



### Class-loader modes

There are two possible values for a class-loader mode:

#### PARENT\_FIRST

The PARENT\_FIRST class-loader mode causes the class loader to first delegate the loading of classes to its parent class loader before attempting to load the class from its local classpath. This is the default for class-loader policy and for standard JVM class loaders.

#### PARENT\_LAST

The PARENT\_LAST class-loader mode causes the class loader to first attempt to load classes from its local classpath before delegating the class loading to its parent. This policy allows an application class loader to override and provide its own version of a class that exists in the parent class loader.

The following settings determine a class loader's mode:

- If the application class-loader policy of an application server is SINGLE, the application class-loader policy of an application server defines the mode for an application class loader.
- If the application class-loader policy of an application server is MULTIPLE, the class-loader mode of an application defines the mode for an application class loader.
- If the WAR class-loader policy of an application is MODULE, the WAR class-loader policy of a Web module defines the mode for a WAR class loader.

---

## Class loader collection

Use this page to manage class-loader instances on an application server. A class loader determines whether an application class loader or a parent class loader finds and loads Java class files for an application.

To view this administrative console page, click **Servers > Application Servers > server\_name > Classloader**.

## ClassLoader ID

States a string unique to the server identifying the class-loader instance. The product assigns the identifier.

## ClassLoader Mode

Specifies the class-loader mode when the application class-loader policy is SINGLE. PARENT\_FIRST causes the class loader to delegate the loading of classes to its parent class loader before attempting to load the class from its local classpath. PARENT\_LAST causes the class loader to attempt to load classes from its local classpath before delegating the class loading to its parent; this allows an application class loader to override and provide its own version of a class that exists in the parent class loader.

## Class loader settings

Use this page to configure a class loader for applications that reside on an application server.

To view this administrative console page, click **Servers > Application Servers > server\_name > Classloader > class\_loader\_ID**.

### ClassLoader ID

States a string unique to the server identifying the class-loader instance. The product assigns the identifier.

Data type                      String

### ClassLoader Mode

Specifies the class-loader mode when the application class-loader policy is SINGLE. PARENT\_FIRST causes the class loader to delegate the loading of classes to its parent class loader before attempting to load the class from its local classpath. PARENT\_LAST causes the class loader to attempt to load classes from its local classpath before delegating the class loading to its parent; this allows an application class loader to override and provide its own version of a class that exists in the parent class loader.

Data type                      String  
Default                          PARENT\_FIRST

---

## Migrating the class-loader Module Visibility Mode setting

WebSphere Application Server Version 4.0.x had a server-wide configuration setting called **Module Visibility Mode**. For Version 5.0, you use application or WAR class-loader policies instead of module visibility modes. The Version 5.0 policies provide additional flexibility because you can configure applications running in a server for an application class-loader policy of SINGLE or MULTIPLE and for a WAR class-loader policy of APPLICATION or MODULE.

To migrate module visibility modes in your Version 4.0.x applications to their equivalents in Version 5.0, change the settings for your Version 4.0.x applications and modules to the Version 5.0 values shown in the table below.

| Version 4.0.x module visibility mode | Version 5.0 application class-loader policy | Version 5.0 WAR class-loader policy |
|--------------------------------------|---|-------------------------------------|
| Server                               | SINGLE                                      | APPLICATION                         |

|               |          |             |
|---------------|----------|-------------|
| Compatibility | SINGLE   | MODULE      |
| Application   | MULTIPLE | APPLICATION |
| Module*       | MULTIPLE | MODULE      |
| J2EE          | MULTIPLE | MODULE      |

\*There is no exact equivalent for the Version 4.0.x Module mode because it isolated EJB modules within an application.

---

## Class loading: Resources for learning



Use the following links to find relevant supplemental information about class loaders. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.



View links to additional information about:

- Programming model and decisions
- Programming instructions and examples
- Programming specifications



### Programming model and decisions

-  **J2EE Class Loading Demystified**  
([http://www7b.boulder.ibm.com/wsdd/library/techarticles/0112\\_deboer/deboer.html](http://www7b.boulder.ibm.com/wsdd/library/techarticles/0112_deboer/deboer.html))
-  **Understanding J2EE Application Server Class Loading Architectures**  
(<http://www.theserverside.com/resources/article.jsp?l=ClassLoading>)

### Programming instructions and examples

-  **Developing and Deploying Modular J2EE Applications with WebSphere Studio Application Developer and WebSphere Application Server**  
([http://www7b.boulder.ibm.com/wsdd/library/techarticles/0206\\_robinson/robinson.html](http://www7b.boulder.ibm.com/wsdd/library/techarticles/0206_robinson/robinson.html))
-  **IBM WebSphere Application Server Programming** (<http://www.mcgraw-hill.co.uk/html/0072224592.html>)

### Programming specifications

-  **Sun's J2EETM Platform Specification**  
(<http://java.sun.com/j2ee/download.html#platformspec>)
-  **Sun's J2EETM Extension Mechanism Architecture**  
(<http://java.sun.com/j2se/1.4/docs/guide/extensions/spec.html>)





---

## Chapter 13. Using EJB query

The EJB query language is used to specify a query over container-managed entity beans. The language is similar to SQL. An EJB query is independent of the bean's mapping to a persistent store.

An EJB query can be used in three situations:

- To define a finder method of an EJB entity bean.
- To define a select method of an EJB entity bean.
- To dynamically specify a query using the `executeQuery()` dynamic API.

Finder and select queries are specified in the bean's deployment descriptor using the `<ejb-ql>` tag. Queries specified in the deployment descriptor are compiled into SQL during deployment. Dynamic queries require the interface provided by WebSphere Application Server Enterprise.

WebSphere's EJB query language is compliant with the EJB QL defined in Sun's EJB 2.0 specification and has additional capabilities as listed in the topic [Comparison of EJB 2.0 specification and WebSphere Query Language](#).

In your WebSphere application, you can define an EJB query in the following ways:

- **Application Assembly Tool.** When assembling an EJB 2.0 entity bean, specify the `<ejb-ql>` tag for the `finder()` or `select()` method.
- **WebSphere Studio Application Developer.** When defining an entity bean, specify the `<ejb-ql>` tag for the `finder` or `select` method.
- **Dynamic query service.** Add the `executeQuery()` method to your application. The dynamic query API is provided as an Enterprise Extension to WebSphere Application Server.

Before using EJB query, familiarize yourself with query language concepts, starting with the topic, [EJB Query Language](#).

### Usage scenario

See the topic ["Example: EJB queries"](#).

---

## EJB query language

An EJB query is a string that contains the following elements:

- a SELECT clause that specifies the EJBs or values to return;
- a FROM clause that names the bean collections;
- an optional WHERE clause that contains search predicates over the collections;
- an optional GROUP BY and HAVING clause (see [Aggregation functions](#));
- an optional ORDER BY clause that specifies the ordering of the result collection.

The SELECT clause is optional in order to maintain compatibility with WebSphere Application Server Version 4.

Collections of entity beans are identified in EJB queries through the use of their abstract schema name in the query FROM clause.

The elements of EJB query language are discussed in more detail in the following related topics.

## Example: EJB queries

Here is an example EJB schema, followed by a set of example queries:

### DeptBean schema

|                                    |  |
|------------------------------------|--|
| Entity bean name (EJB name)        | DeptEJB (not used in query)  |
| Abstract schema name               | DeptBean   |
| Implementation class               | com.acme.hr.deptBean (not used in query)   |
| Persistent attributes (cmp fields) | <ul style="list-style-type: none"> <li>• deptno - Integer (key)</li> <li>• name - String</li> <li>• budget - BigDecimal</li> </ul> |
| Relationships                      | <ul style="list-style-type: none"> <li>• emps - 1:Many with EmpEJB</li> <li>• mgr - Many:1 with EmpEJB</li> </ul>                  |

### EmpBean schema

|                                    |  |
|------------------------------------|--|
| Entity bean name (EJB name)        | EmpEJB (not used in query)   |
| Abstract schema name               | EmpBean  |
| Implementation class               | com.acme.hr.empBean (not used in query)  |
| Persistent attributes (cmp fields) | <ul style="list-style-type: none"> <li>• empid - Integer (key)</li> <li>• name - String</li> <li>• salary - BigDecimal</li> <li>• bonus - BigDecimal</li> <li>• hireDate - java.sql.Date</li> <li>• birthDate - java.util.Calendar</li> <li>• address - com.acme.hr.Address</li> </ul> |
| Relationships                      | <ul style="list-style-type: none"> <li>• dept - Many:1 with DeptEJB</li> <li>• manages - 1:Many with DeptEJB</li> </ul>  |

Address is a serializable object used as cmp field in EmpBean. The definition of address is as follows:

```
public class com.acme.hr.Address extends Object implements Serializable {
    public String street;
    public String state;
    public String city;
    public Integer zip;
    public double distance (String start_location) { ... };
    public String format ( ) { ... };
}
```

The following query returns all departments:

```
SELECT OBJECT(d) FROM DeptBean d
```

The following query returns departments whose name begins with the letters "Web". Sort the result by name:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d.name LIKE 'Web%' ORDER BY d.name
```

The keywords SELECT and FROM are shown in uppercase in the examples but are case insensitive. If a name used in a query is a reserved word, the name must be enclosed in double quotes to be used in the query. There is a list of reserved words later in this document. Identifiers enclosed in double quotes are case sensitive. This example shows how to use a cmp field that is a reserved word:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d."select" > 5
```

The following query returns all employees who are managed by Bob. This example shows how to navigate relationships using a path expression:

```
SELECT OBJECT (e) FROM EmpBean e WHERE e.dept.mgr.name='Bob'
```

A query can contain a parameter which refers to the corresponding value of the finder or select method. Query parameters are numbered starting with 1:

```
SELECT OBJECT (e) FROM EmpBean e WHERE e.dept.mgr.name= ?1
```

This query shows navigation of a multivalued relationship and returns all departments that have an employee that earns at least 50000 but not more than 90000:

```
SELECT OBJECT(d) FROM DeptBean d, IN (d.emps) AS e  
WHERE e.salary BETWEEN 50000 and 90000
```

There is a join operation implied in this query between each department object and its related collection of employees. If a department has no employees, the department does not appear in the result. If a department has more than one employee that earns more than 50000, that department appears multiple times in the result.

The following query eliminates the duplicate departments:

```
SELECT DISTINCT OBJECT(d) from DeptBean d, IN (d.emps) AS e WHERE e.salary > 50000
```

Find employees whose bonus is more than 40% of their salary:

```
SELECT OBJECT(e) FROM EmpBean e where e.bonus > 0.40 * e.salary
```

Find departments where the sum of salary and bonus of employees in the department exceeds the department budget:

```
SELECT OBJECT(d) FROM DeptBean d where d.budget <  
( SELECT SUM(e.salary+e.bonus) FROM IN(d.emps) AS e )
```

A query can contain DB2 style date-time arithmetic expressions if you use java.sql.\* datatypes as CMP fields and your datastore is DB2. Find all employees who have worked at least 20 years as of January 1st, 2000:

```
SELECT OBJECT(e) FROM EmpBean e where year( '2000-01-01' - e.hireDate ) >= 20
```

If the datastore is not DB2 or if you prefer to use java.util.Calendar as the CMP field, then you can use the java millisecond value in queries. The following query finds all employees born before Jan 1, 1990:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.birthDate < 631180800232
```

Find departments with no employees:

```
SELECT OBJECT(d) from DeptBean d where d.emps IS EMPTY
```

Find all employees whose earn more than Bob:

```
SELECT OBJECT(e) FROM EmpBean e, EmpBean b
WHERE b.name = 'Bob' AND e.salary + e.bonus > b.salary + b.bonus
```

Find the employee with the largest bonus:

```
SELECT OBJECT(e) from EmpBean e WHERE e.bonus =
(SELECT MAX(e1.bonus) from EmpBean e1)
```

The above queries all return EJB objects. A finder method query must always return an EJB Object for the home. A select method query can in addition return CMP fields or other EJB Objects not belonging to the home.

The following would be valid select method queries for EmpBean. Return the manager for each department:

```
SELECT d.mgr FROM DeptBean d
```

Return department 42 manager's name:

```
SELECT d.mgr.name FROM DeptBean d WHERE d.deptno = 42
```

Return the names of employees in department 42:

```
SELECT e.name FROM EmpBean e WHERE e.dept.deptno=42
```

Another way to write the same query is:

```
SELECT e.name from DeptBean d, IN (d.emps) AS e WHERE d.deptno=42
```

Finder and select queries allow only a single CMP field or EJBObject in the SELECT clause.

The dynamic query api allows multiple expressions in the SELECT clause. The following query would be a valid dynamic query, but not a valid select or finder query:

```
SELECT e.name, e.salary+e.bonus as total_pay , object(e), e.dept.mgr
FROM EmpBean e
ORDER BY 2
```

The following dynamic query returns the number of employees in each department:

```
SELECT e.dept.deptno as department_number , count(*) as employee_count
FROM EmpBean e
GROUP BY by e.dept.deptno
ORDER BY 1
```

The dynamic query api allows queries that contain bean or value object methods:

```
SELECT object(e), e.address.format( )
FROM EmpBean e EmpBean e
```

## FROM clause

The FROM clause specifies the collections of objects to which the query is to be applied. Each collection is identified either by an abstract schema name and an identification variable, called a range variable, or by a collection member declaration that identifies a multivalued relationship and an identification variable.

Conceptually, the semantics of the query is to first form a temporary collection of tuples R. Tuples are composed of elements from the collections identified in the FROM clause. Each tuple contains one element from each of the collections in the

FROM clause. All possible combinations are formed subject to the constraints imposed by the collection member declarations. If any schema name identifies a collection for which there are no records in the persistent store, then the temporary collection R will be empty.

**Example: FROM clause**

DeptBean contains records 10, 20 and 30 in the persistent store. EmpBean contains records 1, 2 and 3 that are related to department 10 and records 4, 5 that are related to department 20. Department 30 has no related employees.

```
FROM DeptBean d, EmpBean e
```

This forms a temporary collection R that contains 15 tuples.

```
FROM DeptBean d, DeptBean d1
```

This forms a temporary collection R that contains 9 tuples.

```
FROM DeptBean d, IN (d.emps) AS e
```

This forms a temporary collection R that contains 5 tuples. Department 30 because it contains no employees will not be in R. Department 10 will be contained in R three times and department 20 will be contained in R twice.

After forming the temporary collection the search conditions of the WHERE clause will be applied to R and this will yield a new temporary collection R1. The ORDER BY and SELECT clauses are applied to R1 to yield the final result set.

An identification variable is a variable declared in the FROM clause using the operator IN or the optional AS.

```
FROM DeptBean AS d, IN (d.emps) AS e
```

is equivalent to:

```
FROM DeptBean d, IN (d.emps) e
```

An identification variable that is declared to be an abstract schema name is called a range variable. In the query above "d" is a range variable. An identification variable that is declared to be a multivalued path expression is called a collection member declaration. "d" and "e" in the example above are collection member declarations.

Note that the following path expression is illegal as a collection member declaration because it is not multivalued:

```
e.dept.mgr
```

## Inheritance in EJB query

If an EJB inheritance hierarchy has been defined for an abstract schema, using the abstract schema name in a query statement implies the collection of objects for that abstract schema as well as all subtypes.

**Example: Inheritance**

Suppose that bean ManagerBean is defined as a subtype of EmpBean and ExecutiveBean is a subtype of ManagerBean in an EJB inheritance hierarchy. The following query returns employees as well as managers and executives:

```
SELECT OBJECT(e) FROM EmpBean e
```

## Path expressions

An identification variable followed by the navigation operator ( . ) and a cmp or relationship name is a path expression.

A path expression that leads to a cmr field can be further navigated if the cmr field is single-valued. If the path expression leads to a multi-valued relationship, then the path expression is terminal and cannot be further navigated. If the path expression leads to a cmp field whose type is a value object, it is possible to navigate to attributes of the value object.

### Example: Value object

Assume that address is a cmp field for EmpBean, which is a value object.

```
SELECT object(e) FROM EmpBean e
WHERE e.address.distance('San Jose') < 10 and e.address.zip = 95037
```

It is best to use the composer pattern to map value object attributes to relational columns if you intend to search on value attributes. If you store value objects in serialized format, then each value object must be retrieved from the database and deserialized. Value object methods can only be done in dynamic queries.

A path expression can also navigate to a bean method. The method must be defined on either the remote or local bean interface. Methods can only be used in dynamic queries. You cannot mix both remote and local methods in a single query statement.

If the query contains remote methods, the dynamic query must be executed using the query remote interface. Using the query remote interface causes the query service to activate beans and create instances of the remote bean interface

Likewise, a query statement with local bean methods must be executed with the query local interface. This causes the query service to activate beans and local interface instances.

Do not use get methods to access cmp and cmr fields of a bean.

If a method has overloaded definitions, the overloaded methods must have different number of parameters.

Methods must have non-void return types and method arguments and return types must be either primitive types byte, short, int, long, float, double, boolean, char or wrapper types from the following list:

Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Character, java.util.Calendar, java.sql.Date, java.sql.Time, java.sql.Timestamp, java.util.Date

If any input argument to a method is NULL, it is assumed the method returns a NULL value and the method is not invoked.

A collection valued path expression can be used in the FROM clause as a collection member declaration, and with the IS EMPTY, MEMBER OF, and EXISTS predicates in the WHERE clause.

|  |    |
|--|----|
| FROM EmpBean e WHERE e.dept.mgr.name='Bob' | OK |
|--|----|

|  |  |
|--|--|
| FROM EmpBean e WHERE e.dept.emps.name='BOB'                | INVALID — cannot navigate through emps because it is multivalued |
| FROM EmpBean e, IN (e.dept.emps) e1<br>WHERE e1.name='BOB' | OK   |
| FROM EmpBean e WHERE e.dept.emps IS EMPTY                  | OK   |

## WHERE clause

The WHERE clause contains search conditions composed of the following:

- literal values
- input parameters
- expressions
- basic predicates
- quantified predicates
- BETWEEN predicate
- IN predicate
- LIKE predicate
- NULL predicate
- EMPTY collection predicate
- MEMBER OF predicate
- EXISTS predicate
- IS OF TYPE predicate

If the search condition evaluates to TRUE, the tuple is added to the result set.

### Literals

A string literal is enclosed in single quotes. A single quote that occurs within a string literal is represented by two single quotes; For example: "Tom's". A string literal cannot exceed the maximum length that is supported by the underlying persistent datastore.

A numeric literal can be any of the following:

- an exact value such as 57, -957, +66
- any value supported by Java long
- a decimal literal such as 57.5, -47.02
- an approximate numeric value such as 7E3, -57.4E-2

A decimal or approximate numeric value must be in the range supported by the underlying persistent datastore.

A boolean literal can be the keyword TRUE or FALSE and is case insensitive.

### Input parameters

Input parameters are designated by the question mark followed by a number; For example: ?2

Input parameters are numbered starting at 1 and correspond to the arguments of the finder or select method; therefore, a query must not contain an input parameter that exceeds the number of input arguments.

An input parameter can be a primitive type of byte, short, int, long, float, double, boolean, char or wrapper types of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Char, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp or an EJBObject.

An input parameter must not have a NULL value. To search for the occurrence of a NULL value the NULL predicate should be used.

## Expressions

Conditional expressions can consist of comparison operators and logical operators (AND, OR, NOT).

Arithmetic expressions can be used in comparison expressions and can be composed of arithmetic operations and functions, path expressions that evaluate to a numeric value and numeric literals and numeric input parameters.

String expressions can be used in comparison expressions and can be composed of string functions, path expressions that evaluate to a string value and string literals and string input parameters. A cmp field of type char is handled as if it were a string of length 1.

Boolean expressions can be used with = and <> comparison and can be composed of path expressions that evaluate to a boolean value and TRUE and FALSE keywords and boolean input parameters.

Reference expressions can be used with = and <> comparison and can be composed of path expressions that evaluate to a cmr field, an identification variable and an input parameter whose type is an EJB reference

Four different expression types are supported for working with date-time types. For portability the java.util.Calendar type should be used. DB2 style date, time and timestamp expressions are supported if the datastore is DB2 and the CMP field is of type java.util.Date, java.sql.Date, java.sql.Time or java.sql.Timestamp.

A Calendar type can be compared to another Calendar type, an exact numeric literal or input parameter of type long whose value is the standard Java long millisecond value.

The following query finds all employees born before Jan 1, 1990:  
`SELECT OBJECT(e) FROM EmpBean e WHERE e.birthDate < 631180800232`

Date expressions can be used in comparison expressions and can be composed of operators + - , date duration expressions and date functions, path expressions that evaluate to a date value, string representation of a date and date input parameters.

Time expressions can be used in comparison expressions and can be composed of operators + - , time duration expressions and time functions, path expressions that evaluate to a time value, string representation of time and time input parameters.

Timestamp expressions can be used in comparison expressions and can be composed of operators + - , timestamp duration expressions and timestamp functions, path expressions that evaluate to a timestamp value, string representation of a timestamp and timestamp input parameters.

Standard bracketing ( ) for ordering expression evaluation is supported.



The operators and their precedence order from highest to lowest are:

- Navigation operator ( . )
- Arithmetic operators in precedence order:
  - + - unary
  - \* / multiply, divide
  - + - add, subtract
- Comparison operators: =, >, <, >=, <=, <>(not equal)
- Logical operator NOT
- Logical operator AND
- Logical operator OR

In some datastores, a zero length string value ( " " ) is treated as a null value and affects the results of queries. Some datastores perform division between two integer values using integer arithmetic rules and other datastores use non integer rules. This also can affect the results of queries. For portability, avoid the use of zero length string values and division of integer values in an EJB query.

### Null value semantics

The following describe the semantics of NULL values:

- Comparison or arithmetic operations with an unknown (NULL) value yield an unknown value
- Path expressions that contain NULL evaluate to NULL
- The IS NULL and IS NOT NULL operators can be applied to path expressions and return TRUE or FALSE. Boolean operators AND, OR and NOT use three valued logic.

| AND     | True    | False | Unknown |
|---------|---------|-------|---------|
| True    | True    | False | Unknown |
| False   | False   | False | False   |
| Unknown | Unknown | False | Unknown |

| OR      | True | False   | Unknown |
|---------|------|---------|---------|
| True    | True | True    | True    |
| False   | True | False   | Unknown |
| Unknown | True | Unknown | Unknown |

|         | NOT     |
|---------|---------|
| True    | False   |
| False   | True    |
| Unknown | Unknown |

### Example: Null value semantics

```
select object(e) from EmpBean where e.salary > 10 and e.dept.budget > 100
```

If salary is NULL the evaluation of `e.salary > 10` returns unknown and the employee object is not returned. If the cmr field dept or budget is NULL evaluation of `e.dept.budget > 100` returns unknown and the employee object is not returned.

```
select object(e) from EmpBean where e.dept.budget is null
```

If dept or budget is NULL evaluation of e.dept.budget is null returns TRUE and the employee object is returned.

```
select object(e) from EmpBean e , in (e.dept.emps) e1 where e1.salary > 10
```

If dept is NULL, then the multivalued path expression e.dept.emps results in an empty collection (not a collection that contains a NULL value). An employee with a null dept value will not be returned.

```
select object(e) from EmpBean e where e.dept.emps is empty
```

If dept is NULL the evaluation of the predicate in unknown and the employee object is not returned.

```
select object(e) from EmpBean e , EmpBean e1 where e member of e1.dept.emps
```

If dept is NULL evaluation of the member of predicate returns unknown and the employee is not returned.

### Date time arithmetic and comparisons

DATE, TIME and TIMESTAMP values may be compared with another value of the same type. Comparisons are chronological. Date time values can also be incremented, decremented, and subtracted.

If the datastore is DB2, then DB2 string representation of DATE, TIME and TIMESTAMP types can also be used. A string representation of a date or time can use ISO, USA, EUR or JIS format. A string representation of a timestamp uses ISO format.

| Format | Date format | Date examples        | Time format    | Time examples    |
|--------|-------------|----------------------|----------------|------------------|
| ISO    | yyyy-mm-dd  | 1987-02-24 1987-2-24 | hh.mm.ss       | 13.50.00 13.50   |
| USA    | mm/dd/yyyy  | 2/24/1987            | hh:mm AM or PM | 1:50 pm 02:10 AM |
| EUR    | dd.mm.yyyy  | 24.02.1987 24.2.1987 | hh.mm.ss       | 13.50.00 13.55   |
| JIS    | yyyy-mm-dd  | 1987-02-24           | hh:mm:ss       | 13:50 13:50:05   |

#### Example 1: Date time arithmetic comparisons

```
e.hiredate > '1990-02-24'
```

The timestamp of February 24th, 1990 1:50 pm can be represented as follows:

```
'1990-02-24-13.50.00.000000' or  
'1990-02-24-13.50.00'
```

If the datastore is DB2, DB2 decimal durations can be used in expressions and comparisons. A date duration is a decimal(8,0) number that represents the difference between two dates in the format YYYYMMDD. A time duration is a decimal(6,0) number that represents the difference between two time values as HHMMSS. A timestamp duration is a decimal(20,6) number representing the differences between two timestamp values as YYYYMMDDHHMMSS.ZZZZZZ (ZZZZZZ is the number of microseconds and is to the right of the decimal point).

Two date values (or time values or timestamp values) can be subtracted to yield a duration. If the second operand is greater than the first the duration is a negative decimal number. A duration can be added or subtracted from a datetime value to yield a new datetime value.

#### Example 2: Date time arithmetic comparisons

DATE('3/15/2000') - '12/31/1999' results in a decimal number 215 which is a duration of 0 years, 2 months and 15 days.

Durations are really decimal numbers and can be used in arithmetic expressions and comparisons.

( DATE('3/15/2000') - '12/31/1999' ) + 14 > 215 evaluates to TRUE.

DATE('12/31/1999') + DECIMAL(215,8,0) results in a date value 3/15/2000.

TIME('11:02:26') - '00:32:56' results in a decimal number 102930 which is a time duration of 10 hours, 29 minutes and 30 seconds.

TIME('00:32:56') + DECIMAL(102930,6,0) results in a time value of 11:02:26.

TIME('00:00:59') + DECIMAL(240000,6,0) results in a time value of 00:00:59.

e.hiredate + DECIMAL(500,8,0) > '2000-10-01' means compare the hiredate plus 5 months to the date 10/01/2000.

### Basic predicates

Basic predicates can be of two forms

```
expression-1 comparison-operator expression-2  
expression-3 comparison-operator ( subselect )
```

The subselect must not return more than one value and the subselect can not return a type of an EJB reference. Boolean types and reference types only support = and <> comparisons.

#### Example: Basic predicates

```
d.name='Java Development'  
e.salary > 20000  
e.salary > ( select avg(e.salary) from EmpBean e)
```

### Quantified predicates

A quantified predicate compares a value with a set of values produced by a subselect.

```
expression comparison-operator SOME | ANY | ALL ( subselect )
```

The expression must not evaluate to a reference type.

When SOME or ANY is specified the result of the predicate is as follows:

- TRUE if the comparison is true for at least one value returned by the subselect.
- FALSE if the subselect is empty or if the comparison is false for every value returned by the subselect.
- UNKNOWN if the comparison is not true for all of the values returned by the subselect and at least one of the comparisons is unknown because of a null value.

When ALL is specified the result of the predicate is as follows:

- TRUE if the subselect returns empty or if the comparison is true to every value returned by the subselect.
- FALSE if the comparison is false for at least one value returned by the subselect.
- UNKNOWN if the comparison is not false for all values returned by the subselect and at least one comparison is unknown because of a null value.

## **BETWEEN predicate**

The BETWEEN predicate determines whether a given value lies between two other given values.

```
expression [NOT] BETWEEN expression-2 AND expression-3
```

### **Example: BETWEEN predicate**

```
e.salary BETWEEN 50000 AND 60000
```

is equivalent to:

```
e.salary >= 50000 AND e.salary <= 60000
```

```
e.name NOT BETWEEN 'A' AND 'B'
```

is equivalent to:

```
e.name < 'A' OR e.name > 'B'
```

## **IN predicate**

The IN predicate compares a value to a set of values and can have one of two forms:

```
expression [NOT] IN ( subselect )
```

```
expression [NOT] IN ( value1, value2, .... )
```

ValueN can either be a literal value or an input parameter. The expression can not evaluate to a reference type.

### **Example: IN predicate**

```
e.salary IN ( 10000, 15000 )
```

is equivalent to

```
( e.salary = 10000 OR e.salary = 15000 )
```

```
e.salary IN ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

is equivalent to

```
e.salary = ANY ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

```
e.salary NOT IN ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

is equivalent to

```
e.salary <> ALL ( select e1.salary from EmpBean e1 where e1.dept.deptno = 10)
```

## **LIKE predicate**

The LIKE predicate searches a string value for a certain pattern.

```
string-expression [NOT] LIKE pattern [ ESCAPE escape-character ]
```

The pattern value is a string literal or parameter marker of type string in which the underscore ( `_` ) stands for any single character and percent ( `%` ) stands for any sequence of characters ( including empty sequence ). Any other character stands for itself. The escape character can be used to search for character `_` and `%`. The escape character can be specified as a string literal or an input parameter.

If the string-expression is null, then the result is unknown.

If both string-expression and pattern are empty, then the result is true.

### **Example: LIKE predicate**

- `'' LIKE ''` is true

- '' LIKE '%' is true
- e.name LIKE '12%3' is true for '123' '12993' and false for '1234'
- e.name LIKE 's\_me' is true for 'some' and 'same', false for 'soome'
- e.name LIKE '/\_foo' escape '/' is true for '/\_foo', false for 'afoo'
- e.name LIKE '//\_foo' escape '/' is true for '/afoo' and for '/bfoo'
- e.name LIKE '///\_foo' escape '/' is true for '/\_foo' but false for '/afoo'

### NULL predicate

The NULL predicate tests for null values.

single-valued-path-expression IS [NOT] NULL

#### Example: NULL predicate

```
e.name IS NULL
e.dept.name IS NOT NULL
e.dept IS NOT NULL
```

### EMPTY collection predicate

To test if a multivalued relationship is empty, use the following syntax:

collection-valued-path-expression IS [NOT] EMPTY

#### Example: Empty collection predicate

To find all departments with no employees:

```
SELECT OBJECT(d) FROM DeptBean d WHERE d.emps IS EMPTY
```

### MEMBER OF predicate

This expression tests whether the object reference specified by the single valued path expression or input parameter is a member of the designated collection. If the collection valued path expression designates an empty collection the value of the MEMBER OF expression is FALSE.

{single-valued-path-expression | input\_parameter} [NOT] MEMBER [OF] collection-valued-path-expression

#### Example: MEMBER OF predicate

Find employees that are not members of a given department number:

```
SELECT OBJECT(e) FROM EmpBean e , DeptBean d
WHERE e NOT MEMBER OF d.emps AND d.deptno = ?1
```

Find employees whose manager is a member of a given department number:

```
SELECT OBJECT(e) FROM EmpBean e, DeptBean d
WHERE e.dept.mgr MEMBER OF d.emps and d.deptno=?1
```

### EXISTS predicate

The exists predicate tests for the presence or absence of a condition specified by a subselect.

```
EXISTS ( subselect )
EXISTS collection-valued-path-expression
```

The result of EXISTS is true if the subselect returns at least one value or the path expression evaluates to a nonempty collection, otherwise the result is false.

To negate an EXISTS predicate, precede it with the logical operator NOT.

#### Example: EXISTS predicate

Return departments that have at least one employee earning more than 1000000:

```
SELECT OBJECT(d) FROM DeptBean d
WHERE EXISTS ( SELECT 1 FROM IN (d.emps) e WHERE e.salary > 1000000 )
```

Return departments that have no employees:

```
SELECT OBJECT(d) FROM DeptBean d
WHERE NOT EXISTS ( SELECT 1 FROM IN (d.emps) e)
```

The above query can also be written as follows:

```
SELECT OBJECT(d) FROM DeptBean d WHERE NOT EXISTS d.emps
```

### **IS OF TYPE predicate**

The IS OF TYPE predicate is used to test the type of an EJB reference. It is similar in function to the Java instance of operator. IS OF TYPE is used when several abstract beans have been grouped into an EJB inheritance hierarchy. The type names specified in the predicate are the bean abstract names. The ONLY option can be used to specify that the reference must be exactly this type and not a subtype.

```
identification-variable IS OF TYPE ( [ONLY] type-1, [ONLY] type-2, ..... )
```

#### **Example: IS OF TYPE predicate**

Suppose that bean ManagerBean is defined as a subtype of EmpBean and ExecutiveBean is a subtype of ManagerBean in an EJB inheritance hierarchy.

The following query returns employees as well as managers and executives:

```
SELECT OBJECT(e) FROM EmpBean e
```

If you are interested in objects which are employees and not managers and not executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE( ONLY EmpBean )
```

If you are interested in object which are managers or executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE( ManagerBean)
```

The above query is equivalent to the following query:

```
SELECT OBJECT(e) FROM ManagerBean e
```

If you are interested in managers only and not executives:

```
SELECT OBJECT(e) FROM EmpBean e WHERE e IS OF TYPE( ONLY ManagerBean)
```

or:

```
SELECT OBJECT(e) FROM ManagerBean e
WHERE e IS OF TYPE (ONLY ManagerBean)
```

## **Scalar functions**

EJB query contains scalar built-in functions for doing type conversions, string manipulation, and for manipulating date-time values. The list of scalar functions is documented in the topic EJB query: Scalar functions.

#### **Example: Scalar functions**

Find employees hired in 1999:

```
SELECT OBJECT(e) FROM EmpBean e where YEAR(e.hireDate) = 1999
```

The only scalar functions that are guaranteed to be portable across backend datastore vendors are the following:

- ABS
- SQRT
- CONCAT
- LENGTH
- LOCATE
- SUBSTRING

The other scalar functions should be used only when DB2 is the backend datastore.

### **EJB query: Scalar functions**

EJB query contains scalar built-in functions, as listed below, for doing type conversions, string manipulation, and for manipulating date-time values.

#### **Numeric functions**

ABS ( < any numeric datatype > ) -> < any numeric datatype >

SQRT ( < any numeric datatype > ) -> Double

#### **Type conversion functions**

CHAR ( < any numeric datatype > ) -> string

CHAR ( < string > ) -> string

CHAR ( < any datetime datatype > [, Keyword k ] ) -> string

Datetime datatype is converted to its string representation in a format specified by the keyword k. The valid keywords values are ISO, USA, EUR or JIS. If k is not specified the default is ISO.

BIGINT ( < any numeric datatype > ) -> Long

BIGINT ( < string > ) -> Long

The following function converts the argument to an integer n by truncation and returns the date that is n-1 days after January 1, 0001:

DATE ( < date string > ) -> Date

DATE ( < any numeric datatype> ) -> Date

The following function returns date portion of a timestamp:

DATE( timestamp ) -> Date

DATE ( < timestamp-string > ) -> Date

The following function converts number to decimal with optional precision p and scale s.

DECIMAL ( < any numeric datatype > [, p [ ,s ] ] ) -> Decimal

The following function converts string to decimal with optional precision p and scale s.

DECIMAL ( < string > [ , p [ , s ] ] ) -> Decimal

DOUBLE ( < any numeric datatype > ) -> Double

DOUBLE ( < string > ) -> Double

FLOAT ( < any numeric datatype > ) -> Double

FLOAT ( < string > ) -> Double

Float is a synonym for DOUBLE.

INTEGER ( < any numeric datatype > ) -> Integer

INTEGER ( < string > ) -> Integer

REAL ( < any numeric datatype > ) -> Float  
 SMALLINT ( < any numeric datatype > ) -> Short  
 SMALLINT ( < string > ) -> Short  
 TIME ( < time > ) -> Time  
 TIME ( < time-string > ) -> Time  
 TIME ( < timestamp > ) -> Time  
 TIME ( < timestamp-string > ) -> Time  
 TIMESTAMP ( < timestamp > ) -> Timestamp  
 TIMESTAMP ( < timestamp-string > ) -> Timestamp

### String functions

CONCAT ( <string>, <string> ) -> String

The following function returns a character string representing absolute value of the argument not including its sign or decimal point. For example, digits( -42.35) is "4235".

DIGITS ( Decimal d ) -> String

The following function returns the length of the argument in bytes. If the argument is a numeric or datetime type, it returns the length of internal representation.

LENGTH ( < string > ) -> Integer

The following function returns a copy of the argument string where all upper case characters have been converted to lower case.

LCASE ( < string > ) -> String

The following function returns the starting position of the first occurrence of argument 1 inside argument 2 with optional start position. If not found, it returns 0.

LOCATE ( String s1 , String s2 [, Integer start ] ) -> Integer

The following function returns a substring of s beginning at character m and containing n characters. If n is omitted, the substring contains the remainder of string s. The result string is padded with blanks if needed to make a string of length n.

SUBSTRING ( String s , Integer m [ , Integer n ] ) -> String

The following function returns a copy of the argument string where all lower case characters have been converted to upper case.

UCASE ( < string > ) -> String

### Date - time functions

The following function returns the day portion of its argument. For a duration, the return value can be -99 to 99.

DAY ( Date ) -> Integer  
 DAY ( < date-string > ) -> Integer  
 DAY ( < date-duration > ) -> Integer  
 DAY ( Timestamp ) -> Integer  
 DAY ( < timestamp-string > ) -> Integer  
 DAY ( < timestamp-duration > ) -> Integer

The following function returns one more than number of days from January 1, 0001 to its argument.



```
DAYS ( Date ) -> Integer
DAYS ( < Date-string > ) -> Integer
DAYS ( Timestamp ) -> Integer
DAYS ( < timestamp-string > ) -> Integer
```

The following function returns the hour part of its argument. For a duration, the return value can be -99 to 99.

```
HOUR ( Time ) -> Integer
HOUR ( < time-string > ) -> Integer
HOUR ( < time-duration > ) -> Integer
HOUR ( Timestamp ) -> Integer
HOUR ( < timestamp-string > ) -> Integer
HOUR ( < timestamp-duration > ) -> Integer
```

The following function returns the microsecond part of its argument.

```
MICROSECOND ( Timestamp ) -> Integer
MICROSECOND ( < timestamp-string > ) -> Integer
MICROSECOND ( < timestamp-duration > ) -> Integer
```

The following function returns the minute part of its argument. For a duration, the return value can be -99 to 99.

```
MINUTE ( Time ) -> Integer
MINUTE ( < time-string > ) -> Integer
MINUTE ( < time-duration > ) -> Integer
MINUTE ( Timestamp ) -> Integer
MINUTE ( < timestamp-string > ) -> Integer
MINUTE ( < timestamp-duration > ) -> Integer
```

The following function returns the month portion of its argument. For a duration, the return value can be -99 to 99.

```
MONTH ( Date ) -> Integer
MONTH ( < date-string > ) -> Integer
MONTH ( < date-duration > ) -> Integer
MONTH ( Timestamp ) -> Integer
MONTH ( < timestamp-string > ) -> Integer
MONTH ( < timestamp-duration > ) -> Integer
```

The following function returns the second part of its argument. For a duration, the return value can be -99 to 99.

```
SECOND ( Time ) -> Integer
SECOND ( < time-string > ) -> Integer
SECOND ( < time-duration > ) -> Integer
SECOND ( Timestamp ) -> Integer
SECOND ( < timestamp-string > ) -> Integer
SECOND ( < timestamp-duration > ) -> Integer
```

The following function returns the year portion of its argument. For a duration, the return value can be -9999 to 9999.

```
YEAR ( Date ) -> Integer
YEAR ( < date-string > ) -> Integer
YEAR ( < date-duration > ) -> Integer
YEAR ( Timestamp ) -> Integer
YEAR ( < timestamp-string > ) -> Integer
YEAR ( < timestamp-duration > ) -> Integer
```

## Aggregation functions

Queries that return aggregate values can only be used with the dynamic query interface available in WebSphere Application Server Enterprise. However, aggregation functions can be used in non-dynamic queries if the aggregation function is used in a subselect or HAVING clause.

Aggregation functions operate on a set of values to return a single scalar value. The following is an example of an aggregation:

```
SELECT SUM (e.salary + e.bonus) FROM EmpBean e WHERE e.dept.deptno =20
```

This computes the total salary and bonus for department 20.

The aggregation functions are avg, count, max, min and sum. The syntax of an aggregation function is as follows:

```
aggregation-function ( [ ALL | DISTINCT ] expression )
```

or:

```
COUNT( * )
```

The DISTINCT option eliminates duplicate values before applying the function. ALL is the default and does not eliminate duplicates. Null values are ignored in computing the aggregate function except for COUNT(\*) which returns a count of all elements in the set.

MAX and MIN can apply to any numeric, string or datetime datatype and returns the same datatype. SUM and AVG take a numeric type as input. SUM and AVG return numeric type. The actual numeric type returned by SUM and AVG depends on the underlying datastore. COUNT can take any datatype and returns an integer.

If you are using an Informix datastore, the argument to COUNT must be an asterisk or a single valued path expression. The argument to SUM, AVG, MIN, or MAX used with DISTINCT must be a single valued path expression.

The set of values that is used for the aggregate function is determined by the collection that results from the FROM and WHERE clause of the subquery. This set can be divided into groups and the aggregation function applied to each group. This is done by using a GROUP BY clause in the subquery. The GROUP BY clause defines grouping members which is a list of path expressions. Each path expression specifies a field that is a primitive type of byte, short, int, long, float, double, boolean, char, or a wrapper type of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Character, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time or java.sql.Timestamp.

Finder or select queries can not return aggregation function values. In other words, aggregation functions can not appear in the top level SELECT of a finder or select query but can be used in subqueries.

#### **Example: Aggregation functions**

```
SELECT e.dept.deptno, AVG ( e.salary) FROM EmpBean e GROUP BY e.dept.deptno
```

The above query computes the average salary for each department.

In dividing a set into groups, a NULL value is considered equal to another NULL value.

Just as the WHERE clause filters tuples from the FROM clause, the groups can be filtered using a HAVING clause that tests group properties involving aggregate functions or grouping members:

```
SELECT e.dept.deptno, AVG ( e.salary) FROM EmpBean e
GROUP BY e.dept.deptno
HAVING COUNT(*) > 3 AND e.dept.deptno > 5
```

This query returns average salary for departments that have more than 3 employees and the department number is greater than 5.

It is possible to have a HAVING clause without a GROUP BY clause in which case the entire set is treated as a single group to which the HAVING clause is applied.

## SELECT clause

For finder and select queries, the syntax of the SELECT clause is as follows:

```
SELECT [ ALL | DISTINCT ]  
{ single-valued-path-expression | OBJECT ( identification-variable ) }
```

The SELECT clause consists of either a single identification variable that is defined in the FROM clause or a single valued path expression that evaluates to a object reference or CMP value. The keyword DISTINCT can be used to eliminate duplicate references.

For a query that defines a finder method the query must return an object type consistent with the home for which the finder method associated with the query. In other words, a finder method for a department home can not return employee objects.

For dynamic queries the syntax is as follows:

```
SELECT { ALL | DISTINCT } [ selection , ]* selection  
selection ::= { expression [[AS] id ] | scalar-subselect }
```

A scalar-subselect is a subselect that returns a single value.

### Example: SELECT clause

Find all employees that earn more than John:

```
SELECT OBJECT(e) FROM EmpBean ej, EmpBean e  
WHERE ej.name = 'John' and e.salary > ej.salary
```

Find all departments that have one or more employees who earn less than 20000:

```
SELECT DISTINCT e.dept FROM EmpBean e where e.salary < 20000
```

A select method query can have a path expression that evaluates to an arbitrary value:

```
SELECT e.dept.name FROM EmpBean e where e.salary < 2000
```

The above query returns a collection of name values for those departments having employees earning less than 20000.

### Example: Valid dynamic queries

The following are examples of dynamic queries:

```
SELECT e.name, e.salary+e.bonus as total_pay from EmpBean e  
SELECT SUM( e.salary+e.bonus) from EmpBean e where e.dept.deptno = ?1
```

## ORDER BY clause

The ORDER BY clause specifies an ordering of the objects in the result collection:

```
ORDER BY [ order_element , ]* order_element  
order_element ::= { path-expression | integer } [ ASC | DESC ]
```

The path expression must specify a single valued field that is a primitive type of byte, short, int, long, float, double, char or a wrapper type of Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Character, java.util.Calendar, java.util.Date, java.sql.Date, java.sql.Time, java.sql.Timestamp.

ASC specifies ascending order and is the default. DESC specifies descending order.

Integer refers to a selection expression in the SELECT clause.

#### **Example: ORDER BY clause**

Return department objects in decreasing deptno order:

```
SELECT OBJECT(d) FROM DeptBean d ORDER BY d.deptno DESC
```

Return employee objects sorted by department number and name:

```
SELECT OBJECT(e) FROM EmpBean e ORDER BY e.dept.deptno ASC, e.name DESC
```

The following is a valid dynamic query:

```
SELECT OBJECT(e), e.salary+e.bonus as total_pay FROM EmpBean e ORDER BY 2 DESC
```

## **Subqueries**

A subquery can be used in quantified predicates, EXISTS predicate or IN predicate. A subquery should only specify a single element in the SELECT clause. When a path expression appears in a subquery, the identification variable of the path expression must be defined either in the subquery, in one of the containing subqueries, or in the outer query. A scalar subquery is a subquery that returns one value. A scalar subquery can be used in a basic predicate and in the SELECT clause of a dynamic query.

#### **Example: Subqueries**

```
SELECT OBJECT(e) FROM EmpBean e
WHERE e.salary > ( SELECT AVG(e1.salary) FROM EmpBean e1)
```

The above query returns employees who earn more than average salary of all employees.

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.salary >
( SELECT AVG(e1.salary) FROM IN (e.dept.emps) e1 )
```

The above query returns employees who earn more than average salary of their department.

```
SELECT OBJECT(e) FROM EmpBean e WHERE e.salary =
( SELECT MAX(e1.salary) FROM IN (e.dept.emps) e1 )
```

The above query returns employees who earn the most in their department.

```
SELECT OBJECT(e) FROM EmpBean e
WHERE e.salary > ( SELECT AVG(e.salary) FROM EmpBean e1
WHERE YEAR(e1.hireDate) = YEAR(e.hireDate) )
```

The above query returns employees who earn more than the average of employees hired in same year.

## **EJB query restrictions**

An EJB query is compiled into an SQL query and executed against the underlying datastore based on schema mapping of the abstract bean to the datastore schema. The semantics of comparison and arithmetic operations are that of the underlying

datastore. In the case of SQL, note that two strings are equal if the shorter string padded with blanks equals the longer string. For example, 'A' is equal to 'A '. This differs from the equality of strings in the Java language. Arithmetic overflow operations are an error in SQL.

A cmp field can not be used in comparison operations or predicates (except for LIKE) if that cmp field is mapped to a long varchar or large objects (LOB) column or any other column type for which the database server does not support predicates or comparison operations.

A cmp field of any type can be used in a SELECT clause. Fields that can be used in predicates, grouping, or ordering operations must be of the types listed below:

- Primitive types : byte, short, int, long, float, double, boolean, char
- Object types: Byte, Short, Integer, Long, Float, Double, BigDecimal, String, Boolean, Char, java.util.Calendar , java.util.Date
- JDBC types: java.sql.Date, java.sql.Time, java.sql.Timestamp

The field must be mapped to a table column that is compatible in type either by using a "top-down" default mapping generated by the WebSphere deploy tool, or using a "meet-in-the-middle" mapping between compatible types.

In order to search on attributes of a cmp field that is a user-defined value object, you should use a "meet-in-the-middle" mapping and use a composer to map each attribute to a compatible column. The default "top-down" mapping stores the object as a serialized object in a column of type blob, which does not allow searching.

If a cmp field is mapped to a column using a "meet-in-the-middle" mapping with a converter, that field can only be used with the NULL predicate or with basic predicates of the following form:

```
path-expression <comparison> literal_value
path-expression <comparison> input_parameter
```

In this situation, the converter method toData( ) is called to convert the right-hand side of the predicate to an SQL value.

Example of allowable predicate on a cmp field with user defined converter:

```
e.name = 'Chris'
e.name > ?1
e.name IS NULL
```

Examples of unallowable predicates:

```
substring( e.name, 1, 3 ) = 'ABC'
e.salary > d.budget
```

A converter should preserve equality, collating sequence and null values when doing a conversion. Otherwise cmp fields created by the converter should not be used in WHERE, GROUP, HAVING or ORDER clauses of a query.

## EJB Query: Reserved words

The following words are reserved in WebSphere EJB query:

all, as, distinct, empty, false, from, group, having, in, is, like, select, true, union, where

Avoid using identifiers that start with underscore (for example, `_integer`) as these are also reserved.

## EJB query: BNF syntax

```
EJB QL ::= [select_clause] from_clause [where_clause]
          [group_by_clause] [having_clause] [order_by_clause]
from_clause ::= FROM identification_variable_declaration
              [,identification_variable_declaration]*
identification_variable_declaration ::= collection_member_declaration |
                                       range_variable_declaration
collection_member_declaration ::=
    IN ( collection_valued_path_expression ) [AS] identifier
range_variable_declaration ::= abstract_schema_name [AS] identifier
single_valued_path_expression ::=
    {single_valued_navigation | identification_variable}.
    ( cmp_field | method | cmp_field.value_object_attribute |
      cmp_field.value_object_method )
    | single_valued_navigation
single_valued_navigation ::=
    identification_variable.[ single_valued_cmr_field. ]*
    single_valued_cmr_field
collection_valued_path_expression ::=
    identification_variable.[ single_valued_cmr_field. ]*
    collection_valued_cmr_field
select_clause ::= SELECT { ALL | DISTINCT } {
    single_valued_path_expression | identification_variable |
    OBJECT ( identification_variable) }
select_clause_eex ::= SELECT { ALL | DISTINCT } [ selection , ]* selection
selection ::= { expression [[AS] id ] | subselect }
order_by_clause ::=
    ORDER BY [ {single_valued_path_expression | integer} [ASC|DESC],]*
    {single_valued_path_expression | integer}[ASC|DESC]
where_clause ::= WHERE conditional_expression
conditional_expression ::= conditional_term |
                          conditional_expression OR conditional_term
conditional_term ::= conditional_factor |
                    conditional_term AND conditional_factor
conditional_factor ::= [NOT] conditional_primary
conditional_primary ::= simple_cond_expression | (conditional_expression)
simple_cond_expression ::= comparison_expression | between_expression |
    like_expression | in_expression | null_comparison_expression |
    empty_collection_comparison_expression | quantified_expression |
    exists_expression | is_of_type_expression | collection_member_expression
between_expression ::= expression [NOT] BETWEEN expression AND expression
in_expression ::= single_valued_path_expression [NOT] IN
    { (subselect) | ( atom ,]* atom) }
atom = { string-literal | numeric-constant | input-parameter }
like_expression ::= expression [NOT] LIKE
    {string_literal | input_parameter}
    [ESCAPE {string_literal | input_parameter}]
null_comparison_expression ::=
    single_valued_path_expression IS [ NOT ] NULL
empty_collection_comparison_expression ::=
    collection_valued_path_expression IS [NOT] EMPTY
```

```

collection_member_expression ::=
    { single_valued_path_expression | input_paramter } [ NOT ] MEMBER [ OF ]
    collection_valued_path_expression
quantified_expression ::=
    expression comparison_operator { SOME | ANY | ALL } (subselect)
exists_expression ::= EXISTS { collection_valued_path_expression | (subselect) }
subselect ::=
    SELECT [{ ALL | DISTINCT }] expression from_clause [where_clause]
    [group_by_clause] [having_clause]
group_by_clause ::= GROUP BY [single_valued_path_expression,]*
    single_valued_path_expression
having_clause ::= HAVING conditional_expression
is_of_type_expression ::= identifier IS OF TYPE
    ([[ONLY] abstract_schema_name,]* [ONLY] abstract_schema_name)
comparison_expression ::=
    expression comparison_operator { expression | ( subquery ) }
comparison_operator ::= = | > | >= | < | <= | <>
method ::= method_name( [[expression ,]* expression ] )
expression ::= term | expression {+|-} term
term ::= factor | term {*/} factor
factor ::= {+|-} primary
primary ::= single_valued_path_expression | literal |
    ( expression ) | input_parameter | functions
functions ::=
    ABS(expression) |
    AVG([ALL|DISTINCT] expression) |
    BIGINT(expression) |
    CHAR({expression [, {ISO|USA|EUR|JIS}] } ) |
    CONCAT (expression , expression ) |
    COUNT({[ALL|DISTINCT] expression | *}) |
    DATE(expression) |
    DAY({expression } |
    DAYS( expression ) |
    DECIMAL( expression [,integer[,integer]])
    DIGITS( expression ) |
    DOUBLE( expression ) |
    FLOAT( expression ) |
    HOUR ( expression ) |
    INTEGER( expression ) |
    LCASE ( expression ) |
    LENGTH(expression) |
    LOCATE( expression, expression [, expression] ) |
    MAX([ALL|DISTINCT] expression) |
    MICROSECOND( expression ) |
    MIN([ALL|DISTINCT] expression) |
    MINUTE ( expression ) |
    MONTH( expression ) |
    REAL( expression ) |
    SECOND( expression ) |
    SMALLINT( expression ) |
    SQRT ( expression ) |
    SUBSTRING( expression, expression[, expression] ) |
    SUM([ALL|DISTINCT] expression) |
    TIME( expression ) |
    TIMESTAMP( expression ) |
    UCASE ( expression ) |
    YEAR( expression )

```

## Comparison of EJB 2.0 specification and WebSphere query language

| Item   | EJB 2.0 specification | WebSphere Query | WebSphere Enterprise (Dynamic) Query |
|--|-----------------------|-----------------|--------------------------------------|
| Bean methods   | no                    | no              | yes                                  |
| Calendar comparisons                                   | yes                   | yes             | yes                                  |
| Delimited identifiers                                  | no                    | yes             | yes                                  |
| Dependent Value attributes                             | no                    | yes             | yes                                  |
| Dependent Value methods                                | no                    | no              | yes                                  |
| Dynamic Query APIs                                     | no                    | no              | yes                                  |
| EXISTS predicate                                       | no                    | yes             | yes                                  |
| Inheritance  | no                    | yes             | yes                                  |
| Multiple element select clauses                        | no                    | no              | yes                                  |
| Order by   | no                    | yes             | yes                                  |
| Scalar functions                                       | yes *                 | yes             | yes                                  |
| Select clause  | required              | optional        | optional                             |
| SQL Date/time expressions                              | no                    | yes             | yes                                  |
| String comparisons                                     | = and <> only         | = <> > <        | = <> > <                             |
| Subqueries, aggregations, group by, and having clauses | no                    | yes             | yes                                  |

\* EJB 2.0 defines the following scalar functions: abs, sqrt, concat, length, locate and substring. WebSphere query and dynamic query support additional scalar functions as listed in the topic, EJB query: Scalar functions.

---

## Using the dynamic query service

### Before you begin

Consider using the dynamic query service (available with WebSphere Application Server Enterprise) when any of the following are true:

- You do not know the query search criteria until application runtime.
- You need to return multiple cmp or cmr fields from a query (deployment queries allow only a single element to be specified in the SELECT clause). For more information, see the Example: EJB queries article.
- You want to perform aggregation in the query (deployment queries do not allow use of aggregation function SUM, AVG, COUNT, MAX, MIN in the top level SELECT of a query).
- You want to use value object methods or bean methods in the query statement. For more information, see Path expressions.
- You want to interactively test an EJB query during development but do not want to repeatedly deploy your application each time you update a finder or select query.



If you have a query that has a high frequency of execution you should define it as a finder or select method and consider using SQLJ as a deployment option for best performance. The dynamic query service always uses JDBC and must parse and process the EJB query at runtime.

If you need security control over which queries a user can execute, you need to define the queries as finder or select methods and use EJB method authorization. The dynamic query service does not have fine grain security control at this time. You can control who is permitted access to the remote query bean and the local query bean, but once authorized a user can execute any valid query and return any data in the server.

The dynamic query API is a stateless session bean. Using the dynamic query API is similar to using any other J2EE EJB application bean.

The default JNDI name for the query bean is as follows, but your system administrator can change this name:

```
com/ibm/websphere/ejbquery/Query
```

The system administrator might need to install the query.ear application into the application server. The WebSphere product install does this only for the default server.

**Note:** The query.ear file is located in the <WAS\_HOME>/installableApps directory where <WAS\_HOME> is the location of the WebSphere Application Server.

The query bean has both a remote and a local interface to support both remote and local clients.

- remote interface = com.ibm.websphere.ejbquery.Query
- remote home interface = com.ibm.websphere.ejbquery.QueryHome
- local interface = com.ibm.websphere.ejbquery.QueryLocal
- local home interface = com.ibm.websphere.ejbquery.QueryLocalHome

**5.0+** To use the local interface or bean methods on the remote interface of the query bean, you must configure your server to use the following:

```
Application Classloader Policy = SINGLE
```

Using a value of MULTI may result in your application being unable to find the local interface for the query bean home.

The following Jar files comprise the query service:

| Jar           | Location               | Usage                        |
|---------------|------------------------|------------------------------|
| query.jar     | server - AppServer/lib | query parser                 |
| qjcup.jar     | server - AppServer/lib | auxillary classes for parser |
| querybean.jar | server - installedApps | query session bean           |
| querymd.jar   | server - AppServer/lib | auxillary classes for parser |
| queryws.jar   | server - AppServer/lib | adapter classes for runtime  |
| qryclient.jar | client                 | client stubs and classes     |

Steps for this task

1. To execute a query, have your client do a JNDI lookup for the QueryHome and create an instance of the query bean. The query bean contains the executeQuery() method, which takes as parameters the query statement in the form of a string, and input parameters in the form of an array of java.lang.Object values. Remote clients also pass as arguments the size of the result set to return.

The results of the query are returned for remote clients as:

```
com.ibm.websphere.ejbquery.QueryIterator
```

or for local clients as:

```
com.ibm.websphere.ejbquery.QueryLocalIterator
```

If you want to return remote EJB references from the query, or if the query statement contains remote methods, you must use the query remote interface.

If you want to return local EJB references from the query, or if the query statement contains local methods, you must use the query local interface.

Calling the next() method on the iterator returns an instance of com.ibm.websphere.ejbquery.IQueryTuple, which contains the actual data values or object references. The iterator also contains the following methods:

```
getFieldName( int i)  
getFieldsCount( );
```

2. Compile and run your client program with the file **qryclient.jar** in the classpath.

For more details, see the following articles:

- Javadoc ([../javadoc/ee/com/ibm/websphere/ejbquery/Query.html](#)) for package com.ibm.websphere.ejbquery.
- Class IQueryTuple ([../javadoc/ee/com/ibm/websphere/ejbquery/IQueryTuple.html](#))
- Class QueryIterator ([../javadoc/ee/com/ibm/websphere/ejbquery/QueryIterator.html](#))

The WebSphere Sample Gallery contains a complete working example of the dynamic query service.

**Security Considerations.** WebSphere does not have security access control for CMP and CMR fields. You must therefore secure the query bean methods executeQuery(), prepareQuery(), and executePlan(), and the create() method on the query home; otherwise, any user is able to perform a dynamic query and retrieve data from your application.

## Example: Dynamic query remote client

The following is an example of the Dynamic query remote client with added explanation.

```
import com.ibm.websphere.ejbquery.QueryHome;  
import com.ibm.websphere.ejbquery.Query;  
import com.ibm.websphere.ejbquery.QueryIterator;  
import com.ibm.websphere.ejbquery.IQueryTuple;  
import com.ibm.websphere.ejbquery.QueryException;  
  
try {  
String query =  
"select e.name as name , object(e) as emp from EmpBean e where e.salary < 50000";
```

```

InitialContext ic = new InitialContext();

Object obj = ic.lookup("com/ibm/websphere/ejbquery/Query");

QueryHome qh =
    ( QueryHome) javax.rmi.PortableRemoteObject.narrow( obj, QueryHome.class );
Query qb = qh.create();

```

You must specify a maximum size of the query result set. In this example 99 is used.

```

QueryIterator it = qb.executeQuery(query, null, null ,0, 99 );

```

The iterator contains a collection of IQueryTuple objects. Each tuple contains one value of name and one value of object(e) from the query. To display the contents of the query result, use the following code:

```

while (it.hasNext() ) {
    IQueryTuple tuple = (IQueryTuple) it.next();
    System.out.print( it.getFieldName(1) );
    String s = (String) tuple.getObject(1);
    System.out.println( s);

    System.out.println( it.getFieldName(2) );
    Emp e =
        ( Emp) javax.rmi.PortableRemoteObject.narrow( tuple.getObject(2), Emp.class );
    System.out.println( e.getPrimaryKey().toString());
}

```

The output from the program might look something like the following:

```

name Bob
emp 1001
name Dave
emp 298003
...

```

Finally, catch and process any exceptions. An exception might occur because of a syntax error in the query statement or a run-time processing error. The following is an example that catches and processes these exceptions:

```

} catch (QueryException qe) {
    System.out.println("Query Exception "+ qe.getMessage() );
}

```

## Example: Dynamic query from local client

Establish a transaction context prior to calling `executeQuery` so that a call to `begin()` and `commit()` is needed. The `QueryLocalIterator` becomes invalid at end of a transaction. Therefore, you must use the iterator in the same transaction scope as the `executeQuery` call.

```

import com.ibm.websphere.ejbquery.QueryLocalHome;
import com.ibm.websphere.ejbquery.QueryLocal;
import com.ibm.websphere.ejbquery.QueryLocalIterator;
import com.ibm.websphere.ejbquery.IQueryTuple;
import com.ibm.websphere.ejbquery.QueryException;

String query =
    "select e.name, object(e) from EmpBean e where e.salary < 50000 ";

InitialContext ic = new InitialContext();
    QueryLocalHome qh = ( LocalQueryHome) ic.lookup( "java:comp/env/ejb/query" );
QueryLocal qb = qh.create();
userTransaction.begin();
QueryLocalIterator it = qb.executeQuery(query, null, null);
while (it.hasNext() ) {

```

```

IQueryTuple tuple = (IQueryTuple) it.next();
System.out.print( it.getFieldName(1) );
String s = (String) tuple.getObject(1);
System.out.println( s);

System.out.println( it.getFieldName(2) );
EmpLocal e = ( EmpLocal ) tuple.getObject(2);
System.out.println( e.getPrimaryKey().toString());
}
userTransaction.commit();

```

## Dynamic query service performance considerations

You can use the dynamic query service to build and execute queries against entity beans constructed dynamically at runtime, rather than defining them at deployment time. By using dynamic query you gain the flexibility of queries defined at runtime and utilize the power of EJB-Query Language (QL). Apart from supporting all of the capabilities of an EJB-QL query, dynamic query adds additional functionality not available to standard static query. Two examples are the ability to select multiple data fields directly from the bean itself (static queries currently only allow one) and executing business methods directly in the query.

When used as a direct replacement for an equivalent static query, dynamic query is slower than the static variation. This slowdown is due to the need for parsing and building a plan for the query, in addition to executing it. In the static variation, these costs are paid at deploy time. In WebSphere Application Server Enterprise 5.0, expect to see 60% degradation in performance between executing a predefined static query and the equivalent dynamic query. Through optimizations implemented in WebSphere Application Server Enterprise 5.01, a cache for recently used query plans and other performance improvements, this degradation is reduced to about 25%.

You can effectively create more efficient and less resource intensive applications with dynamic query. For example, two data fields are required from the results of a query. Because a standard EJB-QL query can only select one data field, it is necessary to select the entire EJB object and extract the needed data from the returned results through data access methods, possibly traversing Container Managed Relationships (CMR) boundaries in the process. However, when using dynamic query, you can get both pieces of data directly from the query without additional CMR traversal or accessor methods. This principle is the key to evaluating whether or not dynamic query can be used for performance gain. You should review the amount of data that must be retrieved, in addition to the amount of business logic needed to retrieve it, for example, CMR traversal or accessor methods.

Using parameters in the query rather than literal values is another performance consideration. Under most circumstances, it is better to define conditional values as parameters in the query and then pass those parameters through the appropriate mechanisms. By using this method, you have a greater chance of matching a cached query plan and you eliminate the need to parse and build the plan from scratch. For example, "SELECT Object(o) FROM schemaname AS o WHERE o.fieldname LIKE foo", is more appropriately expressed as "SELECT Object(o) FROM schemaname AS o WHERE o.fieldname LIKE ?1" with the value *foo* passed as a parameter to the executeQuery method. The result is that any subsequent execution of a dynamic query structure that is the same, except for different string literal conditions, are registered as a plan cache hit affecting observed performance.

The dynamic query service is inherently handicapped from a performance viewpoint because of the necessary work of parsing the supplied queries dynamically and building the subsequent query plans. Despite this, the added functionality gained through the use of dynamic query, specifically the ability to select multiple data fields in a single query even across CMRs, creates opportunities to utilize dynamic query for the sake of performance improvement. The additional benefits of dynamic query in WebSphere Application Server Enterprise 5.01, like the use of parameters within the query, further improves the performance through caching and reusing a more general set of plans.



---

## Chapter 14. Using the internationalization service

The internationalization service adds APIs and tooling that enable J2EE applications to manage the distribution of internationalization information, or *internationalization context*, necessary to perform localizations within server-side application components. This topic summarizes the steps involved in using the internationalization service.

### Steps for this task

1. If you have an application that uses the WebSphere Version 4.0 internationalization service, review the topic Migrating internationalized applications.
2. Use the internationalization context API within application components to obtain or manage internationalization context.

Servlet and enterprise bean business methods can use internationalization context to perform locale- and time zone-sensitive localizations. EJB client applications, and server components configured to manage internationalization context must use the internationalization context API to set the context elements scoped to their invocations.

#### 5.0.2

You use the internationalization context API within Web service-enabled J2EE client programs and stateless session beans in the same manner as you would in conventional J2EE components, with one exception:

Internationalization context propagated over Web services requests contains a time zone ID, whereas conventional RMI/IIOP requests propagate complete time zone information, including the raw offset, DST information, and so on.

3. Assemble internationalized applications.  
Use the Application Assembly Tool to configure the internationalization type and any container internationalization attributes for the servlets and enterprise beans of your application.

Internationalization type specifies the internationalization policy applicable to a servlet or an enterprise bean and, in particular, indicates whether the application component or its hosting J2EE container will manage internationalization context. Container internationalization attributes can be specified for container-managed servlet and enterprise bean business methods. These attributes tailor a policy by indicating which context the container will scope to an invocation. Configuring internationalization policies declaratively prescribes, by means of the application's deployment descriptor, the distribution and management of context throughout an application.

#### 5.0.2

You configure internationalization type and container internationalization attributes for Web service-enabled stateless session beans in the same manner as you do for conventional beans.

4. Manage the internationalization service.  
Use the administrative console to enable the service on all application servers. By default, the service is enabled within J2EE client environments but is disabled on application servers. You must enable the service on all application servers hosting your application's servlets and enterprise beans in order to use internationalization context.

**5.0.2** This also applies to J2EE Web service client environments and Web service-enabled enterprise beans.

5. Troubleshoot the internationalization service as needed.  
Use the administrative console to enable the trace service to log internationalization service messages when debugging your applications.

---

## Internationalization

An application that can present information to users according to regional cultural conventions is said to be *internationalized*: The application can be configured to interact with users from different localities in culturally appropriate ways. In an internationalized application, a user in one region sees error messages, output, and interface elements in the requested language. Date and time formats, as well as currencies, are presented appropriately for users in the specified region. A user in another region sees output in the conventional language or format for that region.

Historically, the creation of internationalized applications has been restricted to large corporations writing complex systems. Internationalization techniques have traditionally been expensive and difficult to implement, so they have been applied only to major development efforts. However, given the rise in distributed computing and in the use of the World Wide Web, application developers have been pressured to internationalize a much wider variety of applications. This requires making internationalization techniques much more accessible to application developers.

In an application that is not internationalized, the interface that the user sees is unalterably written into the application code. On the other hand, localizing the displayed strings adds a layer of abstraction into the design of the application. Instead of simply printing an error message, an internationalized application represents the error message with some language-neutral information; in the simplest case, each error condition corresponds to a key. To print a usable error message, then, the application looks up the key in the configured message catalog. Each message catalog is a list of keys with associated strings. Different message catalogs provide strings for the different languages supported. The application looks up the key in the appropriate catalog, retrieves the corresponding error message in the requested language, and prints this string for the user.

Localization of text can be used for far more than translating error messages. For example, by using keys to represent each element in a graphical user interface (GUI) and by providing the appropriate message catalogs, the GUI itself (buttons, menus, and so on) can support multiple languages. Extending support to additional languages requires that you provide message catalogs for those languages; in many cases, the application itself needs no further modification.

Internationalization of an application is driven by two variables, the time zone and the locale. The time zone indicates how to compute the local time as an offset from a standard time like Greenwich Mean Time. The locale is a collection of information about language, currency, and the conventions for presenting information like dates. In a localized application, the locale also indicates the message catalog from which an application is to retrieve message strings. A time zone can cover many locales, and a single locale can span time zones. With both time zone and locale, the date, time, currency, and language for users in a specific region can be determined.



WebSphere Application Server Enterprise provides an Internationalization Service, which manages the distribution of locale and time-zone information, or *internationalization context*, within Java applications.

---

## Internationalization service: Overview

In a distributed client-server environment, application processes can run on different machines, configured to different locales, corresponding to different cultural conventions; they can also be located across geographical boundaries. For an understanding of how these differences impact application development, review the topic Internationalization challenges in distributed applications.

The J2EE platform provides support for application components executing on computers with differing endian architecture and code sets, but does not provide dedicated support for application components that run on computers having different locales or time zones.

The conventional method for solving locale and time zone mismatch across remote application components is to pass one or more extra parameters on all business methods needed to convey the client-side locale or time zone to the server. Although simple, this technique has the following limitations when used in EJB applications:

- It is intrusive because it requires that one or more parameters be added to all bean methods in the call chain to locale-sensitive or time zone-sensitive methods.
- It is inherently error-prone.
- It is impracticable within applications that do not afford modification, such as legacy applications.

## The internationalization service solution

The WebSphere internationalization service addresses the challenges posed by locale and time zone mismatch without incurring the limitations of conventional techniques. It does this by systematically managing the distribution of internationalization contexts across the various components of EJB applications, including client applications, enterprise beans, and servlets.

The service works by associating an internationalization context with every thread of execution within an application. When a client-side component invokes a business method, the internationalization service interposes by obtaining the internationalization context associated with the current thread of the client-side process and attaching that context to the outgoing request. On the server-side, the internationalization service again interposes by detaching the context from the incoming request and associating it with the thread of the server-side process on which the business method will execute, effectively scoping the context to the business method. The service propagates internationalization context on subsequent business method invocations in the same manner and thus distributes the context of the originating request over the entire chain of business method invocations.

This basic operation of scoping and propagation is defined precisely by internationalization context management policies. Every application component has a default policy, which can be overridden and tailored for servlets and enterprise beans at development time using WebSphere's Application Assembly Tool (AAT). Internationalization policies specify whether an application component or its hosting J2EE container are to manage internationalization context. For

container-managed components, the policy indicates which internationalization context the container will scope to invocations on that component. Server components configured to manage internationalization context, as well as EJB clients, must use the internationalization context API to manage the internationalization context elements scoped to their invocations.

At execution time, application components can use the internationalization context API to get any element of the internationalization contexts scoped to an invocation. To programmatically access context elements, application components first resolve an internationalization context API reference, then invoke the appropriate API method to access the various context elements, such as the caller locale or the invocation time zone. These elements can be used in calls to Java 2 SDK internationalization API methods; for example, to perform localizations such as formatting messages, configuring dates, or comparing strings.

## **Internationalization challenges in distributed applications**

With the advent of Internet-based business computational models, such as eCommerce, applications increasingly comprise clients and servers that operate in different locales and geographical regions. These differences introduce challenges to the task of designing a sound client-server infrastructure. Specifically, clients and servers can:

- Execute on computers having different endian architectures or code sets.
- Be located in different locales.
- Be located in different time zones.

The following sections describe these three challenges in more detail.

### **Computers with differing endian architectures or code sets**

Clients and servers can reside in computers having different endian architectures: a client could reside in a little-endian CPU, while the server code runs in a big-endian one. As a more complex instance, a client might want to invoke a business method on a server running in a code set different from that of the client.

A client-server infrastructure must define precise endian and code set tracking and conversion rules. Both CORBA and J2EE have addressed the problems of endian and code set mismatches. The language-neutral CORBA formalism uses byte order indicator in all marshalled data streams to indicate the byte order of the originating machine; in case of an endian mismatch, the receiving side can perform byte swapping for endian correction. The code set mismatch is addressed by CORBA using a comprehensive framework for code set conversion.

J2EE has nearly eliminated these problems in a unique way by relying on its Java Virtual Machine (JVM), which encodes all string data in UCS-2 format and externalizes everything in big-endian format. The JVM employs a set of platform-specific programs for interfacing with the native platform. These programs perform any necessary code set conversions between UCS-2 and the native code set of a platform.

### **Computers located in different locales**

Client and server processes can execute in geographical locations having different locales. For example, a Spanish client might invoke a business method upon an object residing on an American server. Some business methods can be locale-sensitive in nature; for example, given a business method that returns a

sorted list of strings, the Spanish client expects that list to be sorted according to the Spanish collating sequence, not in the server's English collating sequence. Since data retrieval and sorting procedures run on the server, the locale of the client must be available in order to perform a legitimate sort.

A similar consideration applies in instances where the server has to return strings containing date, time, currency, exception messages, and so on, formatted according to the client's cultural expectations. Neither the CORBA nor the J2EE specifications have architecturally addressed the locale mismatch problem and other options involving extra parameters are not practical or have limitations. For example, requiring an extra parameter could require interface changes, which is a serious concern for deployed applications.

### **Computers located in different time zones**

Client and server processes can execute in geographical locations having different time zones. To date, all internationalization literature and resources have concentrated mainly on code set and locale-related issues. They have generally ignored the time zone issue, even though business methods can be sensitive to time zone as well as to locale.

For example, suppose that a vendor makes the claim that "orders received before 2:00 PM will be processed by 5:00 PM the same day". The times given, of course, are in the time zone of the server that is processing the order. It is important to know the client's time zone in order to give customers in other time zones the correct times for same-day processing.

Other time zone-sensitive operations include time stamping messages logged to a server, and accessing file or database resources. The concept of Daylight Savings Time (DST) further complicates the time zone issue. Neither the CORBA nor the J2EE specifications address time zone issues adequately, and conventional methods of solving this problem are limited.

---

## **Migrating internationalized applications**

Applications that used the internationalization service in WebSphere Application Server Version 4.0 can use the service in Version 5.0 with no modification. The packaging and structure of the internationalization context API remain identical across releases. Most importantly, the semantics of the API remain as well.

In Version 4.0, the internationalization service did not provide internationalization deployment descriptor policy information to direct how the service manages internationalization context across the various application components. Rather, the service employed the implicit client-side internationalization (CSI) and server-side internationalization (SSI) policies, which dictated how the service managed context according to the type of J2EE container hosting a component. For details, refer to the WebSphere Application Server Version 4.0 Integrated InfoCenter (<http://www.ibm.com/software/webservers/appserv/doc/v40/aee/index.html>). Briefly, all server components in Version 4.0 were SSI, and all EJB client applications were CSI.

In Version 5.0, the internationalization type setting of all server components is configured to "Container" by default. The internationalization service assigns the default container internationalization attribute, "RunAsCaller", to any container-managed (CMI) servlet or enterprise bean invocation lacking a container internationalization attribute. Hence, invocations of server components lacking

internationalization policy information in the deployment descriptor run under the policy, [CMI, RunAsCaller], which is semantically equivalent to the SSI internationalization policy of Version 4.0; EJB client applications run under the logical policy [AMI, RunAsServer], which is equivalent to the CSI policy of Version 4.0.

When migrating a Version 4.0 application to Version 5.0, it is unnecessary to configure the internationalization deployment descriptor information during application assembly because all component invocations execute under semantically equivalent internationalization context management policies.

---

## Assembling internationalized applications

Use the Application Assembly Tool to configure internationalization deployment descriptor information for servlets and enterprise beans.

### Steps for this task

#### 1. (Optional) Set the **internationalization type**.

All servlets and enterprise beans have an internationalization type setting that specifies whether internationalization context is managed by the application component or by its hosting J2EE container during invocations of their respective lifecycle and business methods. Internationalization type can be configured for all server application components except entity beans, which are container-managed only.

By default, all server components employ container-managed internationalization (CMI). The default setting should suffice in most cases; when it does not, modify the internationalization type setting by completing the steps described in one of the following topics:

- Setting the internationalization type for servlets
- Setting the internationalization type for enterprise beans

#### 2. (Optional) Set the **container internationalization attribute**.

You can associate CMI servlets, and business methods of CMI enterprise beans, with a container internationalization attribute that specifies which of three internationalization contexts - **Caller**, **Server**, or **Specified** - the container is to scope to an invocation. When running as specified, the container internationalization attribute also specifies the custom internationalization context elements.

Named container internationalization attributes can be associated with sets of servlets or with sets of EJB business methods. Initially, CMI servlets and business methods implicitly run as caller and do not associate with a container internationalization attribute. When the implicit behavior or an associated attribute setting is unsuitable, configure an attribute by completing the steps described in one of the following topics:

- Configuring container internationalization for servlets
- Configuring container internationalization for enterprise beans

## Setting the internationalization type for servlets

This task sets the internationalization type for a servlet within a Web module.

### Steps for this task

#### 1. Start the Application Assembly Tool.

2. View the servlets of your Web module by selecting *application\_name* > **Web Modules** > *webmodule\_name* > **Web Components** in the left-hand pane.
3. Select your servlet from the left-hand pane.
4. On the **WAS Enterprise** tab on the right, select either **Container** or **Application** from the **Internationalization type** drop-down menu.
5. Click **Apply**.

### Results

The internationalization type setting is assigned to the servlet.

## Configuring container internationalization for servlets

This task configures container internationalization for a servlet within a Web module. Use this procedure to relate one or more servlets to a container internationalization attribute.

### Steps for this task

1. Start the Application Assembly Tool.
2. In the left-hand panel, expand your Web module.
3. Select **Internationalization**.
4. To configure a new attribute, right-click on **Internationalization** and select **New**; otherwise, skip to step 5.
  - a. On the **New Container Internationalization Attribute** panel, enter a **Description** that uniquely identifies the policy.
  - b. Click **Add**. From the **Add Servlets** list, select one or more servlets to which the attribute will apply and click **OK** to exit the panel. The servlet(s) appear in the Web Components list.
  - c. Click **OK**.

The new attribute description is displayed in the **Description** list in the right-hand panel.

5. Select a named **Container Internationalization Attribute** from the Description list.

The fields of the selected attribute are displayed.

6. If desired, re-enter a description (name) that uniquely identifies this attribute. A description is an arbitrary character string.
7. Click the **Add** button adjacent to the **Web Components** list.
8. From the **Add Servlets** list, select one or more servlets to which the attribute applies.
9. Click **OK**.

The servlets appear in the Web Components list. Repeat these steps for any remaining servlets you want configure.

10. Complete the **Run as** field by selecting **Caller**, **Server**, or **Specified**.  
If the servlet is to run as **Specified**, complete the following steps to specify the context elements that the container will scope to servlet invocations; otherwise, click **OK** to exit this panel.
11. (Optional) Enter the **Description** of the specified localization context.  
A description is an arbitrary character string.
12. Complete the **Time Zone ID** fields:

- a. (Optional) Enter the time zone ID **Description**.  
A description is an arbitrary character string.
  - b. Enter the **Time Zone ID**.  
A time zone ID is an arbitrary, non-empty string that identifies a time zone supported by the Java SDK type, `java.util.SimpleTimeZone`. Refer to the topic Container internationalization attributes for details.
13. Complete the **Locales** fields; perform the following steps to construct an ordered list of locales:
- a. Click **Add**.  
The **Add Locales** panel is displayed.
  - b. [Optional] Enter the locale **Description**.  
A description is an arbitrary character string.
  - c. Enter the **Language Code**.  
A language code is an arbitrary string. A valid locale must contain at least a language code or a country code.
  - d. Enter the **Country Code**.  
A country code is an arbitrary string. A valid locale must contain at least a language code or a country code.
- Refer to the section on locales in the topic Container internationalization attributes for details about language codes, country codes, and variants.
14. Click **Apply** to apply your changes and exit this panel.

## Results

The servlets are now configured to run under the associated internationalization settings.

### **Internationalization assembly properties for Web modules**

Use this page to specify which internationalization context the Web container will scope to servlet service method invocations.

By default, a Web container scopes the caller's internationalization context to service method invocations of servlets and Java Server Pages components (JSPs) configured to the **Container** internationalization type. To override this default scoping behavior, create and configure at least one Container Internationalization attribute.

A Container Internationalization attribute associates those servlets (JSPs) that employ container-managed internationalization to one of three internationalization contexts: the Caller context, the Server context, or the context specified in the attribute. For each servlet listed in the attribute, the container scopes the internationalization context that is specified in the attribute's **Run as** field.

Use the following information to configure new and existing Container Internationalization attributes.

**Description:** Describes the attribute.

**Servlet:** Specifies the set of servlets to which the **Run as** field applies.

Click the **Add** button to select the servlets specified by this attribute.

**Run as:** Specifies the internationalization context that the container will use for scoping method invocations of the servlet indicated in the **Servlet** field.

Default      Caller  
Range        Valid values are Caller, Server, and Specified

More information about valid values follows:

### Caller

The container invokes the attribute's servlet with the locales (the accept-languages) of the incoming HTTP request. The container supplies GMT+00:00 for time zone. Select Caller when you want servlets to execute under the invocation context of the client request.

### Server

The container invokes the servlet with the default context elements of the server Java Virtual Machine (JVM). Select Server when you want servlets to execute under the invocation context of the server JVM.

### Specified

The container invokes the attribute's servlet with the context elements specified in the attribute's **Locales** and **Time zone ID** fields. Select Specified when you want servlets to execute under the invocation context elements specified in the attribute. Be sure to complete the **Locales** and **Time zone ID** fields. You can add information about the group of specified context elements in the **Description** field.

**Time zone:** Represents a temporal offset and computes daylight savings information.

Specify the time zone fields within the **Time Zone** panel according to the following instructions:

#### Description

Describes the specified time zone (ID).

**ID**      Specifies a time zone.

Enter a valid time zone ID. A valid ID represents any time zone supported by java.util.TimeZone. Specifically, an ID may be any of the IDs that appear in the list of time zone IDs returned by method java.util.TimeZone.getAvailableIds(), or a custom ID having the form GMT[+|-]hh[:mm]; for example, America/Los\_Angeles and GMT-08:00 are valid time zone IDs.

**Locales:** Specifies an ordered list of locales; a locale represents a specific geographical, cultural, or political region.

#### Description

Describes the specified locale.

#### Language code

Specifies the language spoken within a particular region.

Enter the language code of the new locale. Ideally, language code is one of the lower-case, two-character codes defined by ISO-639. Language code is not restricted to ISO codes and is not a required field; however, a valid locale must contain either a language code or a country code.

### Country code

Specifies the country within a particular region.

Enter the country code of the new locale. Ideally, country code is one of the upper-case, two-character codes defined by ISO-3166. Country code is not restricted to ISO codes and is not a required field; however, a valid locale must specify either a language code or a country code.

### Variant

A vendor-specific code.

Enter the variant of the new locale. **Variant** is not a required field and serves only to supplement the **Language code** and **Country code** fields according to application- or platform-specific requirements.

## Setting internationalization type for enterprise beans

This task sets the internationalization type for an enterprise bean within an EJB module.

**5.0.2** Use this also for a stateless session bean enabled for Web services.

### Steps for this task

1. Start the Application Assembly Tool.
2. View the beans of your EJB module by selecting *application\_name* > **EJB Modules** > *EJBmodule\_name* > **Session Beans** | **Message-Driven Beans** in the left-hand panel.  
Recall that internationalization type cannot be configured on entity beans, which are CMI only.
3. Select your bean from the left-hand panel.
4. On the **WAS Enterprise** tab of the right-hand panel, select either **Container** or **Application** from the **Internationalization type** drop-down menu.
5. Click **Apply**.

### Results

The internationalization type is assigned to the bean.

## Configuring container internationalization for enterprise beans

This task configures container internationalization for enterprise bean business methods. Use this procedure to relate one or more business methods to a container internationalization attribute.

**5.0.2** Use this also for stateless session beans that are enabled for Web services.

### Steps for this task

1. Start the Application Assembly Tool.
2. In the left-hand panel, expand your EJB module.
3. Select **Internationalization**.
4. To configure a new attribute, right-click on **Internationalization** and select **New**; otherwise, skip to step 5.
  - a. On the **New Container Internationalization Attribute** panel, enter a **Description** that uniquely identifies the policy.



- b. Click **Add**. From the **Add Methods** list, select one or more methods to which the attribute will apply and click **OK** to exit the panel. The selections appear in the Methods list.
- c. Click **OK**.

The new attribute description is displayed in the **Description** list.

5. Select a named **Container Internationalization Attribute** from the **Description** list.

The fields of the selected attribute are displayed.

6. If desired, re-enter a description (name) that uniquely identifies this attribute. A description is an arbitrary character string.

7. To specify the list of bean methods to which the attribute applies, click the **Add** button adjacent to the **Methods** list.

The **Add Methods** panel is displayed.

8. Select a bean method to which the attribute applies.
9. Click **Apply**.

The method appears in the methods list. Repeat these steps for each remaining bean method that you want configure.

10. Click **OK** to exit the Add Methods panel.

11. Complete the **Run As** field by selecting **Caller**, **Server**, or **Specified**.

If the bean is to run as **Specified**, complete the following steps to specify the context elements that the container will scope to bean method invocations; otherwise, click **OK** to exit the panel.

12. (Optional) Enter a **Description** of the specified context.

A description is an arbitrary character string.

13. Complete the **Time Zone ID** fields:

- a. (Optional) Enter the **Time zone ID Description**.

A description is an arbitrary character string.

- b. Enter the **Time Zone ID**.

A time zone ID is an arbitrary, non-empty string that identifies a time zone supported by the Java SDK type, `java.util.SimpleTimeZone`. Refer to the topic Container internationalization attributes for details.

14. Complete the **Locales** fields. Complete the following steps to construct an ordered list of locales:

- a. Click **Add**.

The **Add Locales** panel is displayed.

- b. (Optional) Enter the locale **Description**.

A description is an arbitrary character string.

- c. Enter the **Language code**.

A language code is an arbitrary string. A valid locale must contain at least a language code or a country code.

- d. Enter the **Country code**.

A country code is an arbitrary string. A valid locale must contain at least a language code or a country code.

Refer to the topic Container internationalization attributes for details.

15. Click **OK** to apply your changes and to exit this panel.

## Results

The bean methods are now configured to run under the associated internationalization settings.

### **Internationalization assembly settings for EJB modules**

Use this page to specify which internationalization context the EJB container will scope to enterprise bean business method invocations.

By default, the EJB container scopes the caller's internationalization context to business method invocations of enterprise beans configured to the **Container** internationalization type. To override this default scoping behavior, create and configure at least one Container Internationalization attribute.

A Container Internationalization attribute associates business methods of those enterprise beans that employ container-managed internationalization to one of three internationalization contexts: the Caller context, the Server context, or the context specified in the attribute. For each method listed in the attribute, the container scopes the internationalization context that is specified in the attribute's **Run as** field.

Use the information below to configure new and existing Container Internationalization attributes.

**Description:** Describes the attribute.

**Methods:** Specifies the set of EJB methods to which the **Run as** field applies. Click the **Add** button, then select the methods to be specified in the attribute.

**Run as:** Specifies the internationalization context that the EJB container will use for scoping invocations of the business methods indicated in the **Methods** field.

|         |  |
|---------|--|
| Default | Caller   |
| Range   | Valid values are Caller, Server, and Specified |

More information about valid values follows:

#### **Caller**

The container invokes the attribute's methods with the locales and time zone of the incoming client request. For any missing context element, the container supplies the corresponding default element of the server Java Virtual Machine (JVM). Select **Caller** when you want bean methods to execute under the invocation context of the client request.

#### **Server**

The container invokes the attribute's methods with the default context elements of the server JVM. Select **Server** when you want bean methods to execute under the invocation context of the server JVM.

#### **Specified**

The container invokes the attribute's methods with the context elements specified in the attribute's **Locales** and **Time zone ID** fields. Select **Specified** when you want bean methods to execute under the invocation context specified in the attribute. Be sure to complete the **Locales** and **Time zone ID** fields. You can add information about the group of specified context elements in the **Description** field.

**Time zone:** Represents a temporal offset and computes daylight savings information.

**Description**

A description of the specified time zone (ID).

**ID** A short-hand identifier for a time zone.

Enter a valid time zone ID. A valid ID represents any time zone supported by `java.util.TimeZone`. Specifically, an ID can be any of the IDs that appear in the list of time zone IDs returned by method `java.util.TimeZone.getAvailableIds()`, or a custom ID having the form `GMT[+|-]hh[:mm]`; for example, `America/Los_Angeles` and `GMT-08:00` are valid time zone IDs.

**Locales:** Specifies an ordered list of locales; a locale represents a specific geographical, cultural, or political region.

**Description**

Describes the specified locale.

**Language code**

Specifies the language spoken within a particular region.

Enter the language code of the new locale. Ideally, language code is one of the lower-case, two-character codes defined by ISO-639. Language code is not restricted to ISO codes and is not a required field; however, a valid locale must contain either a language code or a country code.

**Country code**

Specifies the country within a particular region.

Enter the country code of the new locale. Ideally, country code is one of the upper-case, two-character codes defined by ISO-3166. Country code is not restricted to ISO codes and is not a required field; however, a valid locale must specify either a language code or a country code.

**Variant**

A vendor-specific code.

Enter the variant of the new locale. **Variant** is not a required field and serves only to supplement the **Language code** and **Country code** fields according to application- or platform-specific requirements.

---

## Using the internationalization context API

Before you begin

Enterprise JavaBeans (EJB) client applications, servlets, and enterprise beans can programmatically obtain and manage internationalization context using the internationalization context API.

The `java.util` and `com.ibm.websphere.i18n.context` packages contain all classes necessary to use the internationalization service within an EJB application. Classes specific to the internationalization service reside in the file `WAS_HOME/lib/i18nctx.jar`. Before compiling application components that import internationalization service classes, add the `i18nctx.jar` file to your `CLASSPATH`.

Steps for this task

1. Gain access to the internationalization context API.

2. Access caller locales and time zone.
3. Access invocation locales and time zone.

### Usage scenario

Each EJB application component uses the internationalization context API differently. Three code examples are provided that illustrate how to use the API within each application type. Differences in API usage, as well as other coding tips, are noted in comments that precede the relevant statement blocks.

**5.0.2** For Web services client applications and stateless session beans, you use the API to obtain and manage internationalization context in the same manner as for EJB clients and enterprise beans, respectively.

## Gaining access to the internationalization context API

This topic describes how to access the internationalization service by resolving a reference to the internationalization context API.

### Tip:

Resolve internationalization context API references once over the lifecycle of an application component, within the initialization method of that component (for example, within the `init()` method of servlets, or within the `SetXxxContext()` method of enterprise beans).

**5.0.2** For Web service client programs, perform the following task to resolve a reference to the internationalization context API during initialization. For stateless session beans enabled for Web services, resolve the reference in the `setSessionContext()` method.

### Steps for this task

1. Resolve a reference to the `UserInternationalization` interface by performing a lookup on the JNDI name `java:comp/websphere/UserInternationalization`.

For example:

```
//-----
// Internationalization context imports.
//-----
import com.ibm.websphere.i18n.context.*;
import javax.naming.*;
...

public class MyApplication {
    ...

    //-----
    // Resolve a reference to the UserInternationalization interface.
    //-----
    InitialContext initCtx = null;
    UserInternationalization userI18n = null;
    final String UserI18nUrl = "java:comp/websphere/UserInternationalization";
    try {
        initCtx = new InitialContext();
        userI18n = (UserInternationalization)initCtx.lookup(UserI18nUrl);
    }
    catch (NamingException ne) {
        // UserInternationalization URL is unavailable.
    }
}
```

If the `UserInternationalization` object is unavailable due to an anomaly or a restriction, the JNDI lookup invocation throws a `javax.naming.NameNotFoundException` that contains the `java.lang.IllegalStateException`.

2. Use the `UserInternationalization` reference to create references to the `CallerInternationalization` or `InvocationInternationalization` objects, which provide access to elements of the Caller or Invocation internationalization contexts, respectively.

The `CallerInternationalization` reference can be bound to the `Internationalization` interface, only; the `InvocationInternationalization` reference can be bound to either the `Internationalization` or the `InvocationInternationalization` interfaces, depending on whether the application requires read-only or read-write access to invocation context. For example:

```
...
//-----
// Resolve references to the Internationalization and
// InvocationInternationalization interfaces.
//-----
Internationalization callerI18n = null;
InvocationInternationalization invocationI18n = null;
try {
    callerI18n = userI18n.getCallerInternationalization();
    invocationI18n = userI18n.getInvocationInternationalization();
}
catch (IllegalStateException ise) {
    // An Internationalization interface(s) is unavailable.
}
```

## Accessing caller locales and time zone

### Before you begin

An application component must first resolve a reference to the `CallerInternationalization` object and then bind it to the `Internationalization` interface.

Every remote invocation of an application component has an associated caller internationalization context associated with the thread running that invocation. Caller context is propagated by the internationalization service and middleware to the target of a request, such as an EJB business method or servlet service method. Perform the following task to access elements of the Caller internationalization context.

**5.0.2** This task also applies to Web service client programs and stateless session beans enabled for Web services.

### Steps for this task

1. Obtain the desired caller context elements.

```
java.util.Locale [] myLocales = null;
try {
    myLocales = callerI18n.getLocales();
}
catch (IllegalStateException ise) {
    // The Caller context is unavailable;
    // is the service started and enabled?
}
...
```

The Internationalization interface contains the following methods to get caller internationalization context elements:

- **Locale [] getLocales()** Returns the list of caller locales associated with the current thread.
- **Locale getLocale()** Returns the first in the list of caller locales associated with the current thread.
- **TimeZone getTimeZone()** Returns the caller SimpleTimeZone associated with the current thread.

The Internationalization interface allows read-only access to internationalization context within application components. Methods of the Internationalization interface are available to all EJB application components and are used in the same manner for each, but the method semantics vary according to the component's type. For instance, when obtaining caller locale within an EJB client application, the interface returns the default locale of the host JVM; in contrast, when obtaining caller context within a servlet service method (for example, doPost() or doGet() methods), the interface returns the first locale (accept-language) propagated within the corresponding HTML request. See Internationalization context for a discussion of how the service propagates internationalization context throughout an application.

2. Use the caller context elements to localize computations under a locale or time zone of the calling process.

```
DateFormat df = DateFormat.getDateInstance(myLocale);
String localizedDate = df.getDateInstance().format(aDateInstance);
...
```

## Accessing invocation locales and time zone

### Before you begin

An application component must first resolve a reference to the InvocationInternationalization object and then bind it to the InvocationInternationalization interface of the internationalization context API.

Every remote invocation of a servlet service or EJB business method has an invocation internationalization context associated with the thread running that invocation. Invocation context is the internationalization context under which servlet and business method implementations execute; it is propagated on subsequent invocations by the internationalization service and middleware. Perform the following task to access elements of the invocation internationalization context.

**5.0.2** This task also applies to Web services clients and stateless session beans enabled for Web services.

### Steps for this task

1. Obtain the desired invocation context elements.

```
java.util.Locale myLocale;
try {
    myLocale = invocationI18n.getLocale();
}
catch (IllegalStateException ise) {
    // The invocation context is unavailable;
    // is the service started and enabled?
}
...
```

The `InvocationInternationalization` interface contains the following methods to both get and set invocation internationalization context elements:

- **Locale [] `getLocales()`**. Returns the list of invocation locales associated with the current thread.
- **Locale `getLocale()`**. Returns the first in the list of invocation locales associated with the current thread.
- **TimeZone `getTimeZone()`**. Returns the invocation `SimpleTimeZone` associated with the current thread.
- **`setLocales(Locale [])`**. Sets the list of invocation locales associated with the current thread to the supplied list.
- **`setLocale(Locale)`**. Sets the list of invocation locales associated with the current thread to a list containing the supplied locale.
- **`setTimeZone(TimeZone)`**. Sets the invocation time zone associated with the current thread to the supplied `SimpleTimeZone`.
- **`setTimeZone(String)`**. Sets invocation time zone associated with the current thread to a `SimpleTimeZone` having the supplied ID.

The `InvocationInternalization` interface allows read and write access to invocation internationalization context within application components. However, according to internationalization context management policies, only components configured to manage internationalization context (AMI components) have write access to invocation internationalization context elements. Calls to set invocation context elements within CMI application components result in a `java.lang.IllegalStateException`. Any differences in how application components can use `InvocationInternationalization` methods are explained in Internationalization context.

2. Use the invocation context elements to localize a computation under a locale or time zone of the invoking process.

```
DateFormat df = DateFormat.getDateInstance(myLocale);
String localizedDate = df.getDateInstance().format(aDateInstance);
...
```

### Usage scenario

In the following code example, locale (`en,GB`) and simple time zone (`GMT`) transparently propagate on the call to the `myBusinessMethod()` method. Server-side application components, such as `myEjb`, can use the `InvocationInternationalization` interface to obtain these context elements.

```
...
//-----
// Set the invocation context under which the business method or
// servlet will execute and propagate on subsequent remote business
// method invocations.
//-----
try {
    invocationI18n.setLocale(new Locale("en", "GB"));
    invocationI18n.setTimeZone(SimpleTimeZone.getTimeZone("GMT"));
}
catch (IllegalStateException ise) {
    // Is the component CMI; is the service started and enabled?
}
myEjb.myBusinessMethod();
```

Within CMI application components, the `Internationalization` and `InvocationInternationalization` interfaces are semantically equivalent, and either of these interfaces can be used to obtain the context associated with the thread on

which that component is running. For instance, both interfaces can be used to obtain the list of locales propagated to the servlet doPost() service method.

## Example: Internationalization context in an EJB client program

The following code example illustrates how to use the internationalization context API within a contained EJB client program.

**5.0.2** This example also applies to Web service client programs.

```
//-----  
// Basic Example: J2EE EJB client.  
//-----  
package examples.basic;  
  
//-----  
// INTERNATIONALIZATION SERVICE: Imports.  
//-----  
import com.ibm.websphere.i18n.context.UserInternationalization;  
import com.ibm.websphere.i18n.context.Internationalization;  
import com.ibm.websphere.i18n.context.InvocationInternationalization;  
  
import javax.naming.InitialContext;  
import javax.naming.Context;  
import javax.naming.NamingException;  
import java.util.Locale;  
import java.util.SimpleTimeZone;  
  
public class EjbClient {  
  
    public static void main(String args[]) {  
  
        //-----  
        // INTERNATIONALIZATION SERVICE: API references.  
        //-----  
        UserInternationalization userI18n = null;  
        Internationalization callerI18n = null;  
        InvocationInternationalization invocationI18n = null;  
  
        //-----  
        // INTERNATIONALIZATION SERVICE: JNDI name.  
        //-----  
        final String UserI18NUrl =  
            "java:comp/websphere/UserInternationalization";  
  
        //-----  
        // INTERNATIONALIZATION SERVICE: Resolve the API.  
        //-----  
        try {  
            Context initialContext = new InitialContext();  
            userI18n = (UserInternationalization)initialContext.lookup(  
                UserI18NUrl);  
            callerI18n = userI18n.getCallerInternationalization();  
            invI18n = userI18n.getInvocationInternationalization ();  
        } catch (NamingException ne) {  
            log("Error: Cannot resolve UserInternationalization: Exception: " + ne);  
        } catch (IllegalStateException ise) {  
            log("Error: UserInternationalization is not available: " + ise);  
        }  
        ...  
  
        //-----  
        // INTERNATIONALIZATION SERVICE: Set invocation context.  
        //  
        // Under Application-managed Internationalization (AMI), contained EJB  
        // client programs may set invocation context elements. The following
```



```

// statements associate the supplied invocation locale and time zone
// with the current thread. Subsequent remote bean method calls will
// propagate these context elements.
//-----
try {
    invocationI18n.setLocale(new Locale("fr", "FR", ""));
    invocationI18n.setTimeZone("ECT");
} catch (IllegalStateException ise) {
    log("An anomaly occurred accessing Invocation context: " + ise );
}
...

//-----
// INTERNATIONALIZATION SERVICE: Get locale and time zone.
//
// Under AMI, contained EJB client programs can get caller and
// invocation context elements associated with the current thread.
// The next four statements return the invocation locale and time zone
// associated above, and the caller locale and time zone associated
// internally by the service. Getting a caller context element within
// a contained client results in the default element of the JVM.
//-----
Locale invocationLocale = null;
SimpleTimeZone invocationTimeZone = null;
Locale callerLocale = null;
SimpleTimeZone callerTimeZone = null;
try {
    invocationLocale = invocationI18n.getLocale();
    invocationTimeZone =
        (SimpleTimeZone)invocationI18n.getTimeZone();
    callerLocale = callerI18n.getLocale();
    callerTimeZone = (SimpleTimeZone)callerI18n.getTimeZone();
} catch (IllegalStateException ise) {
    log("An anomaly occurred accessing I18n context: " + ise );
}

...
} // main

...
void log(String s) {
    System.out.println (((s == null) ? "null" : s));
}
} // EjbClient

```

## Example: Internationalization context in an EJB servlet

The following code example illustrates how to use the internationalization context API within a servlet. Note the `init()` and `doPost()` methods.

```

...
//-----
// INTERNATIONALIZATION SERVICE: Imports.
//-----
import com.ibm.websphere.i18n.context.UserInternationalization;
import com.ibm.websphere.i18n.context.Internationalization;
import com.ibm.websphere.i18n.context.InvocationInternationalization;

import javax.naming.InitialContext;
import javax.naming.Context;
import javax.naming.NamingException;
import java.util.Locale;

public class J2eeServlet extends HttpServlet {

...
//-----

```

```

// INTERNATIONALIZATION SERVICE: API references.
//-----
protected UserInternationalization userI18n = null;
protected Internationalization i18n = null;
protected InvocationInternationalization invI18n = null;

//-----
// INTERNATIONALIZATION SERVICE: JNDI name.
//-----
public static final String UserI18NUrl =
    "java:comp/websphere/UserInternationalization";

protected Locale callerLocale = null;
protected Locale invocationLocale = null;

/**
 * Initialize this servlet.
 * Resolve references to the JNDI initial context and the
 * internationalization context API.
 */
public void init() throws ServletException {

    //-----
    // INTERNATIONALIZATION SERVICE: Resolve API.
    //
    // Under Container-managed Internationalization (CMI), servlets have
    // read-only access to invocation context elements. Attempts to set these
    // elements result in an IllegalStateException.
    //
    // Suggestion: cache all internationalization context API references
    // once, during initialization, and use them throughout the servlet
    // lifecycle.
    //-----
    try {
        Context initialContext = new InitialContext();
        userI18n = (UserInternationalization)initialContext.lookup(UserI18NUrl);
        callerI18n = userI18n.getCallerInternationalization();
        invI18n = userI18n.getInvocationInternationalization();
    } catch (NamingException ne) {
        throw new ServletException("Cannot resolve UserInternationalization" + ne);
    } catch (IllegalStateException ise) {
        throw new ServletException ("Error: UserInternationalization is not
            available: " + ise);
    }
    ...
} // init

/**
 * Process incoming HTTP GET requests.
 * @param request Object that encapsulates the request to the servlet
 * @param response Object that encapsulates the response from the
 * Servlet.
 */
public void doGet(
    HttpServletRequest request,
    HttpServletResponse response)
    throws ServletException, IOException {
    doPost(request, response);
} // doGet

/**
 * Process incoming HTTP POST requests
 * @param request Object that encapsulates the request to
 * the Servlet.
 * @param response Object that encapsulates the response from
 * the Servlet.
 */

```

```

public void doPost(
    HttpServletRequest request,
    HttpServletResponse response)
    throws ServletException, IOException {

    ...
    //-----
    // INTERNATIONALIZATION SERVICE: Get caller context.
    //
    // The Internationalization service extracts the accept-languages
    // propagated in the HTTP request and associates them with the
    // current thread as a list of locales within the caller context.
    // These locales are accessible within HTTP Servlet service methods
    // using the caller internationalization object.
    //
    // If the incoming HTTP request does not contain accept languages,
    // the service associates the server's default locale. The service
    // always associates the GMT time zone.
    //
    //-----
    try {
        callerLocale = callerI18n.getLocale(); // caller locale
        // the following code enables you to get invocation locale,
        // which depends on the Internationalization policies.
        invocationLocale = invI18n.getLocale(); // invocation locale
    } catch (IllegalStateException ise) {
        log("An anomaly occurred accessing Invocation context: " + ise);
    }
    // NOTE: Browsers may propagate accept-languages that contain a
    // language code, but lack a country code, like "fr" to indicate
    // "French as spoken in France." The following code supplies a
    // default country code in such cases.
    if (callerLocale.getCountry().equals(""))
        callerLocale = AccInfoJBean.getCompleteLocale(callerLocale);

    // Use iLocale in JDK locale-sensitive operations, etc.
    ...
} // doPost

...
void log(String s) {
    System.out.println (((s == null) ? "null" : s));
}
} // CLASS J2eeServlet

```

## Example: Internationalization context in a session bean

The following code example illustrates how to perform a localized operation using the internationalization service within a session bean.

**5.0.2** This example also applies to Web service-enabled session beans.

```

...
//-----
// INTERNATIONALIZATION SERVICE: Imports.
//-----
import com.ibm.websphere.i18n.context.UserInternationalization;
import com.ibm.websphere.i18n.context.Internationalization;
import com.ibm.websphere.i18n.context.InvocationInternationalization;

import javax.naming.InitialContext;
import javax.naming.Context;
import javax.naming.NamingException;
import java.util.Locale;

/**

```

```

* This is a stateless Session Bean Class
*/
public class J2EESessionBean implements SessionBean {

    //-----
    // INTERNATIONALIZATION SERVICE: API references.
    //-----
    protected UserInternationalization    userI18n = null;
    protected InvocationInternationalization  invI18n = null;

    //-----
    // INTERNATIONALIZATION SERVICE: JNDI name.
    //-----
    public static final String UserI18NUrl =
        "java:comp/websphere/UserInternationalization";
    ...

    /**
     * Obtain the appropriate internationalization interface
     * reference in this method.
     * @param ctx javax.ejb.SessionContext
     */
    public void setSessionContext(javax.ejb.SessionContext ctx) {

        //-----
        // INTERNATIONALIZATION SERVICE: Resolve the API.
        //-----
        try {
            Context initialContext = new InitialContext();
            userI18n = (UserInternationalization)initialContext.lookup(
                UserI18NUrl);
            invI18n = userI18n.getInvocationInternationalization();
        } catch (NamingException ne) {
            log("Error: Cannot resolve UserInternationalization: Exception: " + ne);

        } catch (IllegalStateException ise) {
            log("Error: UserInternationalization is not available: " + ise);
        }
    } // setSessionContext

    /**
     * Set up resource bundle using I18n Service
     */
    public void setResourceBundle()
    {
        Locale invLocale = null;

        //-----
        // INTERNATIONALIZATION SERVICE: Get invocation context.
        //-----
        try {
            invLocale = invI18n.getLocale();
        } catch (IllegalStateException ise) {
            log ("An anomaly occurred while accessing Invocation context: " + ise );
        }
        try {
            Resources.setResourceBundle(invLocale);
            // Class Resources provides support for retrieving messages from
            // the resource bundle(s). See Currency Exchange sample source code.
        } catch (Exception e) {
            log("Error: Exception occurred while setting resource bundle: " + e);
        }
    } // setResourceBundle

    /**
     * Pass message keys to get the localized texts
     * @return java.lang.String []

```

```

    * @param key java.lang.String []
    */
    public String[] getMsgs(String[] key) {
        setResourceBundle();
        return Resources.getMsgs(key);
    }

    ...
    void log(String s) {
        System.out.println((s == null) ? ";null" : s));
    }
} // CLASS J2EESessionBean

```

## Internationalization context API: Programming reference

Application components programmatically manage internationalization context through the `UserInternationalization`, `Internationalization`, and `InvocationInternationalization` interfaces in the `com.ibm.websphere.i18n.context` package. The following code example introduces the internationalization context API:

```

public interface UserInternationalization {
    public Internationalization getCallerInternationalization();
    public InvocationInternationalization
        getInvocationInternationalization();
}

public interface Internationalization {
    public java.util.Locale[] getLocales();
    public java.util.Locale getLocale();
    public java.util.TimeZone getTimeZone();
}

public interface InvocationInternationalization
    extends Internationalization {
    public void setLocales(java.util.Locale[] locales);
    public void setLocale(java.util.Locale jmLocale);
    public void setTimeZone(java.util.TimeZone timeZone);
    public void setTimeZone(String timeZoneId);
}

```

### UserInternationalization interface

The `UserInternationalization` interface provides factory methods for obtaining references to the `CallerInternationalization` and `InvocationInternationalization` context objects. Use these references to access elements of the caller and invocation contexts correlated to the current thread.

Methods of the `UserInternationalization` interface:

#### Internationalization `getCallerInternationalization()`

Returns a reference implementing the `Internationalization` interface that allows access to elements of the caller internationalization context correlated to the current thread. If the service is disabled, this method throws an `IllegalStateException`.

#### InvocationInternationalization `getInvocationInternationalization()`

Returns a reference implementing the `InvocationInternationalization` interface. If the service is disabled, this method throws an `IllegalStateException`.

### Internationalization interface

The Internationalization interface declares methods affording read-only access to internationalization context. Given a caller or invocation internationalization context object created with the UserInternationalization interface, bind the object to the Internationalization interface in order to get elements of that context type. Observe that caller internationalization context can be accessed only through this interface.

Methods of the Internationalization interface:

**Locale[] getLocales()**

Returns the chain of locales within the internationalization context (object) bound to the interface, provided the chain is not null; otherwise this method returns a chain of length(1) containing the default locale of the JVM.

**Locale getLocale()**

Returns the first in the chain of locales within the internationalization context (object) bound to the interface, provided the chain is not null; otherwise this method returns the default locale of the JVM.

**TimeZone getTimeZone()**

Returns the caller time zone (that is, the SimpleTimeZone) associated with the current thread, provided the time zone is non-null; otherwise this method returns the process time zone.

### InvocationInternationalization interface

The InvocationInternationalization interface declares methods affording read and write access to InvocationInternationalization context. Given an invocation internationalization context object created with the UserInternationalization interface, bind the object to the InvocationInternationalization interface in order to get and set elements of the invocation context.

**Note:** According to the container-managed internationalization (CMI) policy, all set methods, setXxx(), throw an IllegalStateException when called within a CMI servlet or enterprise bean.

Methods of the InvocationInternationalization interface:

**void setLocales(java.util.Locale[] locales)**

Sets the chain of locales to the supplied chain, *locales*, within the invocation internationalization context. The supplied chain can be null or have length(>= 0). When the supplied chain is null or has length(0), the service sets the chain of invocation locales to an array of length(1) containing the default locale of the JVM. Null entries can exist within the supplied locale list, for which the service substitutes the default locale of the JVM on remote invocations.

**void setLocale(java.util.Locale locale)**

Sets the chain of locales within the invocation internationalization context to an array of length(1) containing the supplied locale, *locale*. The supplied locale can be null, in which case the service instead sets the chain to an array of length(1) containing the default locale of the JVM.

**void setTimeZone(java.util.TimeZone timeZone)**

Sets the time zone within the invocation internationalization context to the supplied time zone, *time zone*. If the supplied time zone is not an exact instance of java.util.SimpleTimeZone or is null, the service instead sets the invocation time zone to the default time zone of the JVM.

### **void setTimeZone(String timeZoneId)**

Sets the time zone within the invocation internationalization context to the `java.util.SimpleTimeZone` having the supplied ID, *timeZoneId*. If the supplied time zone ID is null or invalid (that is, it does not appear in the list of IDs returned by the `java.util.TimeZone.getAvailableIds()` method) the service sets the invocation time zone to the simple time zone having an ID of GMT, an offset of 00:00, and otherwise invalid fields.

### **Internationalization context**

An *internationalization context* is a distributable collection of internationalization information containing an ordered list, or chain, of locales and a single time zone, where the locales and time zone are instances of Java SDK types, `java.util.Locale` and `java.util.TimeZone`. A locale chain is ordered according to the user's preference.

The internationalization service manages and makes available two varieties of internationalization context: the *caller context*, which represents the caller's localization environment, and the *invocation context*, which represents the localization environment under which a business method runs. Server application components use elements of the caller and invocation internationalization contexts to appropriately tailor locale-sensitive and time zone-sensitive computations.

**Note:** The internationalization service does not support time zone types other than Java SDK type `java.util.SimpleTimeZone`. Unsupported time zone types silently map to the default time zone of the JVM when supplied to internationalization context API methods. For a complete description of the `java.util.Locale`, `java.util.TimeZone`, and `java.util.SimpleTimeZone` types, refer the Java SDK API documentation.

### **Caller context**

Caller internationalization context contains the locale chain and time zone received on incoming EJB business method and servlet service method invocations; it is the internationalization context propagated from the calling process. Use caller context elements within server application components to localize computations to the calling component. Caller context is read-only and can be accessed by all application components by using the **Internationalization** interface of the internationalization context API.

Caller context is computed in the following manner: On an EJB business method or servlet service method invocation, the internationalization service extracts the internationalization context from the incoming request and scopes this context to the method as the caller context. For any missing or null context element, the service inserts the corresponding default element of the JVM (for example, `java.util.Locale.getDefault()` or `java.util.TimeZone.getDefault()`.)

**5.0.2** The service performs a similar insertion whenever missing or null Caller context elements are encountered on invocations of stateless session beans that are enabled for Web services.

Formally, caller context is the invocation context of the calling business method or application component.

### **Invocation context**

Invocation internationalization context contains the locale chain and time zone under which EJB business methods and servlet service methods execute. It is managed by either the hosting container or the application component, depending on the applicable internationalization policy. On outgoing business method requests, it is the context that propagates to the target process. Use invocation context elements to localize computations under the specified settings of the current application component.

Invocation context is computed in the following manner: On an incoming business method or servlet service method invocation, the internationalization service queries the associated context management policy. If the policy is container-managed internationalization (CMI), the container scopes the context designated by the policy to the invocation; otherwise the policy is application-managed internationalization (AMI), and the container scopes a vacuous context to the invocation that can be altered by the method implementation.

Application components can access invocation context elements through both the **Internationalization** and **InvocationInternationalization** interfaces of the internationalization context API. Invocation context elements can be set (overwritten) under the application-managed internationalization policy only.

On an outgoing business method request, the service obtains the currently scoped invocation context and attaches it to the request. This outgoing exported context becomes the caller context of the target invocation. When supplying invocation context elements, either for export on outgoing requests or through the API, the internationalization service always provides the most recent element set using the API; also, the service supplies the corresponding default element of the JVM for any null invocation context element.

**5.0.2** Because the internationalization context that is propagated over Web services (SOAP) requests contains a time zone ID rather than the entire state of a `java.lang.SimpleTimeZone` object, time zone information might be lost when a Web service-enabled client program or session bean becomes involved in remote business computation.

### **Internationalization context: Propagation and scope**

The scope of internationalization context is implicit. Every EJB client application, servlet service method, and EJB business method invocation has two internationalization contexts under which it executes. For each application component invocation, the container enters the caller context and the invocation context, as indicated by the pertinent internationalization policy, into scope before the container delegates to the actual implementation. When the implementation returns, the service removes these contexts from scope. The internationalization service supplies no programmatic mechanism for components to explicitly manage the scope of internationalization context.

The service scopes internationalization context differently with respect to application component type:

- EJB client programs (contained)
- Servlets
- Enterprise beans
- **5.0.2** Web service client programs (contained)



- **5.0.2** Stateless session beans that are enabled for Web services

Internationalization context observes by-value semantics over remote method requests, meaning that changes to internationalization context elements scoped to an invocation do not affect the corresponding elements of the internationalization context scoped to the remote calling process. Also, modifications to context elements obtained using the internationalization context API do not affect the corresponding elements scoped to the invocation.

### **EJB client programs (contained)**

Before it invokes the `main()` method of a client program, the J2EE client container introduces into scope invocation and caller internationalization contexts containing null elements. These contexts remain in scope throughout the life of program. EJB client programs are the base in a chain of remote business method invocations and, technically, do not have a logical caller context. Accessing a caller context element yields the corresponding default element of the client JVM. On outgoing EJB business method requests, the internationalization service propagates the invocation context to the target process. Any unset (null) invocation context elements are replaced with the default of the JVM when exported by the internationalization context API or by outgoing requests.

#### **Tip:**

To propagate values other than the JVM defaults to remote business methods, EJB client programs, as well as AMI servlets or enterprise beans, must set (override) elements of the invocation context. To learn how to set invocation context elements, see the topic *Accessing invocation locales and time zone*.

### **Servlets**

On every servlet service method (`doGet()`, `doPost()`) invocation, the J2EE web container introduces caller and invocation internationalization contexts into scope before delegating to the service method implementation. The caller context contains the accept-languages propagated in the HTTP servlet request, typically from a Web browser. The invocation context contains whichever context is indicated by the container internationalization attribute of the internationalization policy associated with the servlet. Any unset (null) invocation context elements are replaced with the default of the server JVM when exported by the internationalization context API or by outgoing requests. The caller and invocation contexts remain effective until immediately after the implementation returns, at which time the container removes them from scope.

### **Enterprise beans**

On every EJB business method invocation, the J2EE EJB container introduces caller and invocation internationalization contexts into scope before delegating to the business method implementation. The caller context contains the internationalization context elements imported from the incoming IIOP request; if the incoming request lacks a particular internationalization context element, the container scopes a null element. The invocation context contains whichever context is indicated by the container internationalization attribute of the internationalization policy associated with the business method. On outgoing EJB business method requests, the service propagates the invocation context to the target process. Any unset (null) invocation context elements are replaced with the

default of the server JVM when exported by the internationalization context API or by outgoing requests. The caller and invocation contexts remain effective until immediately after the implementation returns, at which time the container removes them from scope.

Consider a simple WebSphere EJB application having a Java client that invokes remote bean method, `myBeanMethod()`. On the client side, the application could use the Internationalization Service API to set invocation context elements. When the client calls `myBeanMethod()`, the service exports the client's invocation context to the outgoing request. On the server side, the service detaches the imported context from the incoming request and scopes it to the method as its caller context; it also scopes the invocation context to the method as indicated by the associated internationalization context management policy. The EJB container then calls the `myBeanMethod()`, which can use the internationalization context API to access elements of either the caller or invocation contexts. When `myBeanMethod()` returns, the EJB container removes these contexts from scope.

### **Web service client programs (contained) 5.0.2**

Before it calls the `main()` method of a Web service client program, the J2EE client container introduces into scope both invocation and caller internationalization contexts that contain null elements. These contexts remain in scope throughout the duration of program execution. Web service client programs are the base in a chain of remote business method invocations and, technically, do not have a logical caller context. Accessing a Caller context element yields the corresponding default element of the client virtual machine.

On outgoing Web service requests, the internationalization service transparently creates a Simple Object Access Protocol (SOAP) header block that contains the invocation context associated with the current thread; the SOAP representation of invocation context is propagated through the request to the target process. Any unset (that is, null) invocation context elements are replaced with the default element of the JVM when exported by the internationalization context API or by outgoing requests. Also, because the header contains only a time zone ID, the additional state of the invocation context's time zone object (`java.lang.SimpleTimeZone`) might be lost, because it does not get propagated through the request.

#### **Tip:**

To propagate values other than the JVM defaults to remote business methods, Web service client programs, as well as AMI servlets or enterprise beans, must set (override) elements of the invocation context. For more information, see "Accessing invocation locales and time zone."

### **Stateless session beans that are enabled for Web services 5.0.2**

On every method invocation of a Web service-enabled bean, the EJB container introduces caller and invocation internationalization contexts into scope before delegating control to the business method implementation. The caller context contains the internationalization context elements imported from the SOAP header block of the incoming request. If the incoming request lacks a particular internationalization context element, the container introduces a null element into

scope. The invocation context contains whichever context is indicated by the container internationalization attribute of the internationalization policy that is associated with the business method.

On outgoing EJB business method requests, the service propagates the invocation context to the target process. Any unset (that is, null) invocation context elements are replaced with the default element of the server JVM when exported by the internationalization context API or by outgoing requests. The caller and invocation contexts remain effective until immediately after control returns from the business method implementation, at which time the container removes them from scope.

On outgoing Web service requests, the internationalization service transparently creates a SOAP header block that contains the invocation context associated with the current thread. The SOAP representation of the invocation context is propagated through the request to the target process. Any unset (that is, null) invocation context elements are replaced with the default element of the JVM when exported by the internationalization context API or by outgoing requests.

### Thread association considerations

The Web and EJB containers scope internationalization contexts to a method by associating it with the thread that executes the method's implementation. Similarly, methods of the internationalization context API either associate context with, or obtain context associated with, the thread on which these methods execute. In cases where new threads are spawned within an application component (for instance, a user-generated thread inside the `service()` method of a servlet, or a system-generated event handling thread in an AWT client) the internationalization contexts associated with the parent thread does not automatically transfer to the newly-spawned thread. In such instances, the service exports the default locale and time zone of the JVM on any remote business method request and on any API calls executed on the new thread. If the default context is inappropriate, the desired invocation context elements must be explicitly associated to the new thread using the `setXxx()` methods of the **InvocationInternationalization** interface. Currently, internationalization context management policies allow invocation context to be set within EJB client programs, as well as within servlets, session beans, and message-driven beans employing application-managed internationalization.

### Example: Internationalization context in a SOAP header

The following code example illustrates how internationalization context is represented within the Simple Object Access Protocol (SOAP) header of a Web service request.

```
<InternationalizationContext>
  <Locales>
    <Locale>
      <LanguageCode>ja</LanguageCode>
      <CountryCode>JP</CountryCode>
      <VariantCode>Nihonbushi</VariantCode>
    </Locale>
    <Locale>
      <LanguageCode>fr</LanguageCode>
      <CountryCode>FR</CountryCode>
    </Locale>
    <Locale>
      <LanguageCode>en</LanguageCode>
      <CountryCode>US</CountryCode>
    </Locale>
  </Locales>
  <TimeZoneID>JST</TimeZoneID>
</InternationalizationContext>
```

## Internationalization context: Management policies

Internationalization policies declaratively prescribe how J2EE application components or their hosting containers (the service) will manage internationalization context on component invocations. There are two internationalization context management policies applicable to all component types:

- Application-managed internationalization (AMI)
- Container-managed internationalization (CMI)

These policies are represented in two parts:

- Internationalization type
- Container internationalization attribute

The service defines a default, or implicit, internationalization policy for every application component type. At development time, assemblers can override the default policy for server component types by explicitly configuring their internationalization type, and optional container internationalization attributes, using the WebSphere Application Assembly Tool. Policies configured during assembly are preserved in the application's deployment descriptor.

All components have an internationalization type that indicates whether it is AMI or CMI; that is, whether a component is to deploy under the application-managed or the container-managed internationalization policy. Application assemblers can set the internationalization type for servlets, session beans, and message-driven beans. Entity beans are implicitly CMI and EJB clients are implicitly AMI; neither can be configured otherwise.

For CMI servlets and enterprise beans, optional container internationalization attributes can be specified to indicate which invocation internationalization context the container is to scope to service or business methods. A CMI service or business method invocation can run under the context of the caller's process, under the default context of the server's JVM, or under a custom context specified in the attribute. Assemblers can specify one container internationalization attribute per disjoint set of CMI servlets within a Web module, or one Attribute per disjoint set of business methods of CMI beans within an EJB module. In other words, a container internationalization attribute can be associated with more than one method, but a method cannot be associated with more than one attribute.

When a WebSphere application server launches an application, the internationalization service collects policy information from the deployment descriptor, then uses this information to construct and associate an internationalization policy to every component invocation. A policy is denoted as:  
[<Internationalization Type>,<Container Internationalization Attribute>]

There are several cases where the deployment descriptor appears to lack policy information, for example: EJB client applications have no configurable internationalization policy settings; AMI components do not have container internationalization attributes; and you are not required to specify container internationalization attributes for CMI components. When the service cannot obtain the explicit internationalization type and container attribute settings from a well-formed deployment descriptor, it implicitly inserts the appropriate setting into the policy.

The service observes the following conventions when applying policies to invocations:

- Servlets (service) and EJB business methods lacking all internationalization policy information in the deployment descriptor implicitly execute under policy [CMI,RunAsCaller].
- CMI servlets and business methods lacking a container internationalization attribute in the deployment descriptor implicitly execute under policy [CMI,RunAsCaller].
- AMI servlets and business methods always lack container internationalization attributes in the deployment descriptor, but implicitly execute under the logical policy [AMI,RunAsServer].
- EJB clients always lack internationalization policy information in the deployment descriptor. By definition, EJB clients are implicitly AMI and run under the invocation context of the JVM; they execute under the logical policy [AMI,RunAsServer].

For conditions other than these, such as a malformed deployment descriptor, refer to the topic Internationalization service errors.

Internationalization policies for EJB clients and HTTP clients cannot be configured using the Application Assembly Tool; HTTP clients do, however, run under the language priority settings of the hosting Web browser. These settings are configurable under the options dialog of most Web browsers; refer to your Web browser's documentation for details.

### **Internationalization type**

Every server application component has an *internationalization type* setting that indicates whether the invocation internationalization context is to be managed by the component or by the hosting J2EE container.

Server application components can be deployed to use one of two types of internationalization context management:

- Application-managed internationalization (AMI)
- Container-managed internationalization (CMI)

A server component may be deployed as AMI or CMI, but not both; CMI is the default. The setting applies to the entire component, on every invocation. Use the Application Assembly Tool to configure the internationalization type for servlets, session beans, and message-driven beans; entity beans are CMI and cannot be configured otherwise. EJB client applications do not have an internationalization type setting, but are implicitly AMI.

### **Application-managed internationalization (AMI)**

Under the AMI deployment policy, component developers assume complete control over the invocation internationalization context. AMI components can use the internationalization context API to programmatically set invocation context elements.

AMI components are expected to manage invocation context. Invocations of AMI components implicitly run under the default locale and time zone of the hosting JVM. Invocation context elements not set using the API default to the corresponding elements of the JVM when accessed through the API or when exported on business methods. To export context elements other than the JVM defaults, AMI servlets, AMI enterprise beans, and EJB client applications must set (overwrite) invocation elements using the internationalization context API. Moreover, the container logically suspends caller context imported on AMI servlet

lifecycle method and AMI EJB business method invocations. To continue propagating the context of the calling process, AMI servlets and enterprise beans must use the API to transfer caller context elements to the invocation context.

Specify AMI for server components that have internationalization context management requirements not supported by container-managed internationalization (CMI).

### **Container-managed internationalization (CMI)**

CMI is the preferred internationalization context management policy for server application components; it is also the default policy. Under CMI, the internationalization service collaborates with the Web and EJB containers to set the invocation internationalization context for servlets and enterprise beans. The service sets invocation context according to the container internationalization attribute of the policy associated with a servlet (service method) or an EJB business method.

A CMI policy contains a container internationalization attribute that indicates which internationalization context the container is to scope to an invocation. For details, see topic Container internationalization attributes. By default, invocations of CMI components run under the caller's internationalization context; or rather, they adhere to the implicit policy [CMI,RunasCaller] whenever the servlet or business is not associated with an attribute in the deployment descriptor. For complete details, see the topic Internationalization context: Management policies.

Methods within CMI components can obtain elements of the invocation context using the internationalization context API, but cannot set them. Any attempt to set invocation context elements within CMI components results in a `java.lang.IllegalStateException`.

Specify container-managed internationalization for server application components requiring standard internationalization context management, then specify the container internationalization attributes for CMI servlets and for business methods of CMI enterprise beans that should not run under the caller's internationalization context.

### **Container internationalization attributes**

The internationalization policy of every CMI servlet and EJB business method has a *container internationalization attribute* that specifies which internationalization context the container is to scope to its invocation.

The container internationalization attribute has three main fields:

- Run as
- Locales
- Time zone ID

As a convenience, you can create named container internationalization attributes and associate them to the following:

- Subsets of CMI servlets within a Web module
- Subsets of business methods of CMI enterprise beans within an Enterprise JavaBeans (EJB) module

- **5.0.2** Subsets of business methods of Web service-enabled session beans. In the following descriptions, the term *supported enterprise bean* refers to both CMI enterprise beans and Web service-enabled session beans.

### Run-as field

The **Run-as** field specifies one of three types of invocation context that a container can scope to a method. For servlet service and EJB business methods, the container constructs the invocation internationalization context according to the **Run as** field setting and associates this context to the current thread before delegating to the method's implementation.

With the Application Assembly Tool, the **Run as** field is configurable for any CMI servlet and business method of a CMI enterprise bean.

- **5.0.2** You can also configure the **Run as** field for Web service business methods.

By default, invocations of servlet service methods and EJB business methods implicitly run as caller (`RunAsCaller`) unless the **Run as** field of a policy's attribute specifies otherwise. EJB client applications and AMI server components always run as server (`RunAsServer`).

Invocation context types specifiable with the **Run as** field are:

**Caller** The container invokes the method under the internationalization context of the calling process. For any missing context element, the container supplies the corresponding default context element of the JVM. Select run as caller when you want the invocation to execute under the invocation context of the calling process.

**Server** The container invokes the method under the default locale and time zone of the JVM. Select run as server when you want the invocation to execute under the invocation context of the JVM.

### Specified

The container invokes the method under the internationalization context specified in the attribute. Select run as specified when you want the invocation to execute under the custom invocation context specified in the policy, then provide the custom context elements by completing the Locales and Time zone ID fields.

**Note:** JMS messages do not contain internationalization context. Although container-managed message-driven beans can be configured to run as caller, the container associates the default elements of the server process when invoking the `onMessage()` method of any message-driven bean configured as [`CMI`, `RunAsCaller`].

### Locales field

The **Locales** field specifies an ordered list of locales that the container scopes to an invocation. With the Application Assembly Tool, the **Locales** field is configurable for CMI servlets and for business methods of supported enterprise beans that run as specified.

A locale represents a specific geographical, cultural, or political region and contains three fields:

- **Language code.** Ideally, language code is one of the lower-case, two-character codes defined by ISO-639; however, language code is not restricted to ISO codes and is not a required field. A valid locale must specify a language code if it does not specify a country code.
- **Country code.** Ideally, country code is one of the upper-case, two-character codes defined by ISO-3166; however, country code is not restricted to ISO codes and is not a required field. A valid locale must specify a country code if it does not specify a language code.
- **Variant.** Variant is a vendor-specific code. Variant is not a required field and serves only to supplement the language and country code fields according to application- or platform-specific requirements.

A valid locale must specify at least a language code or a country code; the variant is always optional. The first locale of the list is returned when accessing invocation context using the internationalization context API method `getLocale()`.

### Time zone ID field

The **Time zone ID** field specifies a shorthand identifier for a time zone that the container scopes to an invocation. With the Application Assembly Tool, the **Time zone ID** field is configurable for CMI servlets and for CMI EJB business methods that run as specified.

**5.0.2** You can also configure the **Time zone ID** field for Web service business methods.

A time zone represents a temporal offset and computes daylight savings information. A valid ID indicates any time zone supported by the SDK type, `java.util.TimeZone`. Specifically, a valid ID is any of the IDs appearing in the list of time zone IDs returned by method `java.util.TimeZone.getAvailableIds()`, or a custom ID having the form `GMT[+|-]hh[:mm]`; for example, `America/Los_Angeles`, `GMT-08:00` are valid time zone IDs.

---

## Managing the internationalization service

To use internationalization context in an Enterprise JavaBeans (EJB) application, the internationalization service must be enabled in the run-time environments for all server-side components (servlets and enterprise beans) as well as all client-side components (EJB client applications).

During installation of this product, if you installed the **Additional services** option but do not require the internationalization service, disable the service on all J2EE clients and application servers. Disabling the service eliminates any possible performance degradation incurred by the implicit distribution of internationalization resources.

**Note:** The internationalization service cannot be enabled for HTTP clients because support for internationalization in that case is provided by the browser, not by the application server.

### Steps for this task

1. As needed, enable or disable the internationalization service for servlets and enterprise beans.



The service is disabled by default within WebSphere enterprise application servers. You enable the service by using either the administrative console or the wsadmin tool.

**5.0.2** This step applies also to Web service-enabled session beans.

2. As needed, enable or disable the internationalization service for EJB clients within the hosting WebSphere J2EE client environments.

The service is enabled by default within the WebSphere J2EE client container.

**5.0.2** This step applies also to Web service clients.

## Enabling the internationalization service for servlets and enterprise beans

Any servlet or enterprise bean can use internationalization context if the internationalization service is enabled within the hosting WebSphere application server instance.

**5.0.2** This task applies also to Web service-enabled enterprise beans.

### Steps for this task

1. Start the administrative console.
2. Select **Servers > Application Servers > *server\_name* > Internationalization Service**.
3. Enable the internationalization service.
  - a. If not already selected, select the **Startup** checkbox.
  - b. Click **OK**.

### Results

When the **Startup** setting is selected, the application server automatically initializes, starts, and enables the internationalization service whenever the server starts. If you change this setting, be sure to restart the application server in order for the new setting to take effect.

To disable the service, clear the **Startup** checkbox. In this case, the internationalization service is initialized but is neither started nor enabled when the application server starts.

### Usage scenario

Alternatively, the internationalization service can be enabled from the command line by using the wsadmin tool. To do this, start the wsadmin tool and enter the following commands:

```
set x [$AdminConfig list I18NService]
$AdminConfig modify $x { { enable true } }
$AdminConfig save
exit
```

### What to do next

If you enable or disable the internationalization service, be sure to stop and then restart the application server in order for the new setting to take effect.

## Internationalization service settings

Use this page to enable or disable the internationalization service.

The internationalization service manages the implicit propagation and scoping of locale and time zone information, called *internationalization context*, within WebSphere Enterprise applications. When the service is enabled, server-side application components can use the internationalization context API to programmatically manage locale and time zone information, or to use this information with the J2SE Internationalization API to perform server-side localizations.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Internationalization Service** .

**Startup:** Specifies whether the server will attempt to start the internationalization service.

|         |                                      |
|---------|--------------------------------------|
| Default | Selected                             |
| Range   | Valid values are Selected or Cleared |

More information about valid values follows:

### Selected

When the application server starts, it attempts to start the internationalization service automatically.

### Cleared

The server does not try to start the internationalization service. If internationalization is to be used in applications that run on this server, the system administrator must select this property then restart the server.

## Enabling the internationalization service for EJB clients

By default, the internationalization service is enabled for use within Enterprise JavaBeans (EJB) client applications whenever the **i18nctx.jar** file is in the CLASSPATH setting constructed by the launchClient tool.

**5.0.2** The internationalization service is also enabled for Web service-enabled clients.

When invoking a Java client application, the launchClient tool sets the CLASSPATH to include the i18nctx.jar file and then activates the WebSphere J2EE client (container), which initializes, starts, and enables the service before delegating to the specified application.

To disable the service for all application server instances in your installation, remove the i18nctx.jar file from the *install\_root/lib* directory. This prevents the file from inadvertently being included in the CLASSPATH setting constructed by the launchClient tool.

To selectively disable the service, include the argument `-CCDI18NService.enable=false` or `-CCDI18NService.enable=no` when invoking the launchClient tool.

---

## Troubleshooting the internationalization service

To have your application server emit trace statements for the internationalization service, specify the appropriate trace string to the server's diagnostic trace service.

### Steps for this task

1. Start the administrative console.
2. Select **Servers > application servers > *server\_name* > Diagnostic Trace Service**.
3. Select the **Enable Trace** checkbox.
4. In the **Trace Specification** field, type the following as a continuous string (no spaces and no line breaks):

```
com.ibm.ws.i18n.context.*=all=enabled:  
com.ibm.websphere.i18n.context.*=all=enabled
```
5. Click **OK**.
6. Click **Save** on the taskbar.
7. Click **Save** in the **Save to Master Repository** panel.

### Results

These settings enable the internationalization service trace when you start or restart the corresponding application server.

## Internationalization service errors

The following conditions can occur while your internationalized application is running. These conditions might cause the internationalization service not to start, to throw instances of `IllegalStateException`, or to exercise default behaviors:

- The service is disabled
- The service is not started
- Invalid context element
- Missing context element
- Invalid policy
- Missing policy

If you encounter unexpected or exceptional behavior, the problem is likely related to one of these conditions. You need to examine the trace log to investigate these conditions, which requires that you configure the diagnostic trace service to generate messages about internationalization service function. To do this, see the topic [Troubleshooting the internationalization service](#).

### **The service is disabled**

The internationalization service does not initialize and start when the service's startup setting is cleared. The service generates a message indicating whether it is enabled or disabled. Applications cannot access the internationalization API when the service is disabled. If an application attempts a JNDI lookup to obtain the `UserInternationalization` reference, the lookup fails with a `NamingException` indicating the reference could not be found. In addition, the service does not scope (propagate) internationalization context on incoming (outgoing) business method invocations.

### **The service is not started**

The internationalization service is operational whenever it is in the STARTED state. For example, if an application attempts to access internationalization context and the service is not started, the API throws an `IllegalStateException`. In addition, the service does not provide runtime support for servlets and enterprise beans.

As an application server progresses through its lifecycle, it initializes, starts, stops, and terminates (destroys) the internationalization service. If an anomaly occurs during initialization, the service does not start. Once the service has been started, its state can change to BLOCKED in the event that a serious error occurs. The service generates a message for every state change.

If a trace message indicates that the service is not STARTED, examine previous messages to determine the problem. For instance, the internationalization service does not start if the activity service is unavailable and a message is displayed to that effect during initialization of the internationalization service.

During startup, the following messages indicate potential configuration or run-time problems:

**No ORB support**

The service could not obtain an instance of the ORB. This is a fatal error. Examine the logs for information.

**No TCM support**

The service could not obtain an instance of its thread context manager. This is a fatal error. Examine the logs for information.

**No IIOP (Activity service) support**

The service could not register with the Activity service. This is a fatal error. The internationalization service cannot propagate or receive context on IIOP requests without Activity service support. Review the logs for error conditions related to the Activity service.

**No AsynchBeans support**

The service could not register into the AsynchBeans environment. This warning indicates that the AsynchBeans environment cannot support internationalization context. If the application server should have AsynchBeans support, verify that the `asynchbeans.jar` and `asynchbeansimpl.jar` files exist in the classpath and review the trace log for any AsynchBeans error conditions.

**No EJB container support**

The service could not register with the EJB container. This is a warning that the internationalization service cannot support enterprise beans. Without EJB container support, internationalization, contexts do not scope properly to EJB business methods. Review the trace log for any EJB container-related error conditions.

**No Web container support**

The service could not register with the Web container. This is a warning that the internationalization service cannot support servlets and Java Server Pages (JSPs). Without Web container support, internationalization contexts do not scope properly to servlet service methods. Review the trace log for any Web container-related error conditions.

**No Meta-data support**

The service could not register with the meta-data service. This is a warning that the internationalization service cannot process the internationalization policies within application deployment descriptors. Without meta-data

support, the service associates the default internationalization context management policy, [CMI, RunAsCaller], to every servlet lifecycle method and enterprise bean business method invocation. Review the trace log for any meta-data service-related error conditions.

#### **No JNDI (Name service) support**

The service could not bind the UserInternationalization object into the namespace. This is a fatal error. Application components are unable to access internationalization context API references, and are therefore unable to access internationalization context elements. Review the trace log for any Naming (JNDI) service-related error conditions.

#### **No API support**

The service could not obtain an instance of an internationalization context API object. This is a fatal error. Application components are unable to access internationalization context API references, and are therefore unable to access internationalization context elements.

#### **Invalid context element**

The service detected an invalid internationalization context element. For example, the internationalization service does not support TimeZone instances of a type other than java.util.SimpleTimeZone. If the service encounters an invalid element, it logs a message and substitutes the corresponding default element of the JVM.

#### **Missing context element**

The service detected a missing internationalization context element. Incoming requests, for example from application servers not supporting the internationalization service will lack internationalization context. When the service attempts to access a caller internationalization context element, which does not exist in this case, it logs a message and substitutes the corresponding default element of the JVM.

Whenever possible, the internationalization service should be enabled within all clients and hosting application servers comprising a WebSphere enterprise application. For more information see the topic Managing the internationalization service.

#### **Invalid policy**

The internationalization service detected a malformed internationalization policy in the application deployment descriptor. At execution, the service replaces the malformed attribute with the appropriate default. For instance, if the internationalization type for an entity bean is set to **Application** during the execution of a servlet or EJB business method invocation, the service logs the inconsistency and enforces the **Container** setting instead.

Also, AMI application components do have an implicit container internationalization attribute. By default they run as server. The service silently enforces the implicit policy, [AMI, RunAsServer], and logs messages to this effect.

Invalid container internationalization attributes are likely to occur when specifying the Locales and Time zone ID fields. When encountering invalid Locales and Time zone ID within attributes, the service replaces each with the corresponding default element of the JVM. Be sure to follow the guidelines provided in the topic Assembling internationalized applications.

## Missing policy

The service detected a missing internationalization policy. At execution, the service replaces the missing policy with the appropriate default. For instance, if the internationalization type is missing for a servlet or enterprise bean, the service sets the attribute to Container.

Container internationalization attributes are not mandatory for CMI application components. In the event that a CMI servlet or EJB business method lacks a container internationalization attribute, the service silently enforces the implicit policy [CMI, RunAsCaller].

When an application lacks internationalization policies in its deployment descriptor, or meta-data support is unavailable, the service logs a message and applies the policy [CMI, RunAsCaller] on every servlet service method and EJB business method invocation.

For more information, see the following topics:

- Assembling internationalized applications
- Container internationalization attributes
- Internationalization type
- Migrating internationalized applications

## Internationalization service exceptions

The internationalization service employs one exception:

`java.lang.IllegalStateException`. This exception indicates one of the following things:

- An application component attempted an operation not supported by the service's programming model.

`IllegalStateException` is thrown whenever a server application component whose internationalization type is set to container-managed Internationalization (CMI) attempts to set invocation context. This is a violation of the CMI policy, under which servlets and enterprise beans cannot modify their invocation internationalization context.

- An anomaly occurred that disabled the service.

For instance, if the internationalization service does not properly initialize, the JNDI lookup on the `UserInternationalization` URL throws a `javax.naming.NameNotFoundException` containing an instance of `IllegalStateException`. Refer to the trace log to determine the reason for failure and, if necessary, contact your IBM support representative.

---

## Internationalization: Resources for learning

Use the following links to find relevant supplemental information about internationalization. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to this product but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming instructions and examples
- Programming specifications

### Programming instructions and examples

- ✓ **Java internationalization tutorial**  
(<http://java.sun.com/docs/books/tutorial/i18n/index.html>)  
An online tutorial that explains how to use the Java 2 SDK Internationalization API.

### Programming specifications

- 📄 **Java 2 SDK, Standard Edition Documentation: Internationalization**  
(<http://java.sun.com/j2se/1.3/docs/guide/intl/>)  
The Java internationalization documentation from Sun Microsystems, including a list of supported locales and encodings.
- **5.0.2** 📄 **Java Specification Request 150, Internationalization Service for J2EE** (<http://www.jcp.org/en/jsr/detail?id=150>)  
The specification of the J2EE internationalization service that is currently being developed through the Java Community Process.
- **5.0.2** 📄 **W3C, Web Services Internationalization Task Force**  
(<http://www.w3.org/2002/05/i18n-recharter/WG-charter.html#web-services>)  
The task force of the W3C's Internationalization Working Group responsible for investigating the internationalization of Web services, in particular, the dependence of Web services on language, culture, region, and locale-related contexts.
- 📄 **Making the WWW truly World Wide** (<http://www.w3.org/International/>)  
The W3C's effort to make World Wide Web technology work with the many writing systems, languages, and cultural conventions of the global community:
- 📄 **developerWorks - Unicode**  
(<http://www.ibm.com/developerworks/unicode/>)  
Articles on various subjects relating to Unicode, from IBM's developerWorks.





---

## Chapter 15. Application profiling

Application profiling enables you to configure multiple access intent policies on the same entity bean. Application profiling reflects the fact that different invocations against the same entity can require different kinds of support from the server run-time environment. For more information, see the topic [Application Profiling: Overview](#).

### Steps for this task

1. Assembling applications for application profiles.

This topic describes how to configure tasks, create application profiles, and configure tasks on profiles, using the Application Assembly Tool.

2. Managing application profiles.

This topic describes how to add and remove tasks from application profiles using the administrative console.

3. Using the TaskNameManager API.

This topic describes how to programmatically set the current task name, but you should use this technique sparingly. Wherever possible, use the declarative method instead, which results in more portable function.

---

### Application profiling: Overview

Application profiling enables developers to identify particular units of work to the WebSphere Application Server Version 5 run-time environment. The run time can tailor its support to the exact requirements of that unit of work. Access intent is currently the only run time component that makes use of the application profiling functionality. For example, you can configure one transaction to load an entity bean with strong update locks and configure another transaction to load the same entity bean without locks.

Application profiling introduces two new concepts in order to achieve this function: "tasks" and "profiles".

**Tasks** A task is a named unit of work within a distributed application. Unit of work in this case means a unique path within the application that may or may not correspond to a transaction or activity session. The name of the path is typically assigned declaratively to a J2EE client or servlet, or to the method of an enterprise bean. This point of configuration marks the head of a graph or subgraph identified by the name of the task; the task name flows from the head of the graph downstream on all subsequent IIOP requests, identifying each subsequent invocation along the graph as belonging to the developer-configured task.

#### **Profiles**

A profile is simply a set of policies that are configured not only on the components of an application, but on a set of tasks as well. When an invocation on a bean (whether by a finder method, a cmr getter, or a dynamic query) requires data to be retrieved from the back-end system, the current task associated with the request is used to determine the exact requirement of the transaction; the same bean loads and behaves differently in the context of the task to profile mapping. Each profile provides the developer an opportunity to reconfigure the application's

access intent. If a request is operating in the absence of a task, the run-time environment uses the access intent configuration external to the application profiles.

## Tasks

Tasks are named units of work. They are the mechanism by which the run-time environment determines which access intent policies to apply when an entity bean's data is loaded from the back-end system.

Application profiles enable developers to configure an entity bean with multiple access intent policies; if there are  $n$  instances of profiles in a given application, each bean can be configured with as many as  $n$  access intent policies.

A task is a unit of work that is given a name by a developer. A task is assigned to any thread executing within a J2EE component, then propagated implicitly across all IOP requests. The WebSphere Application Server run-time environment queries the task at the invocation of any entity bean, and establishes the appropriate access intent policy with which an entity instance will be associated. A task typically corresponds to the execution of a concrete and high-level job within the application.

If an entity bean is loaded in a unit of work that is not associated with a task, or is associated with a task that is unassociated with an application profile, the method-level access intent configuration is applied. If a unit of work is associated with a task that is configured with an application profile, the bean-level access intent configuration within the appropriate application profile is applied.

For example, consider a school district application that calls through a session bean in order to interact with student records. One method on the session bean allows administrators to modify the students' records; another method supports student requests to view their own records. Without application profiling, the two tasks would operate anonymously and the run-time environment would be unable to distinguish work operating on behalf of one task or the other. To optimize the application, a developer can configure one of the methods on the session bean with the task "updateRecords" and the other method on the session bean with the task "readRecords". When registered with an application profile that has the student bean configured with the appropriate locking access intent, the "updateRecords" task is assured that it is not unnecessarily blocking transactions that need to only read the records.

Developers declare tasks using the Application Assembly Tool. Tasks can be declared in three ways:

- Tasks can be associated with J2EE application clients, servlets, and JSPs; requests from these components are then associated with the appropriate task.
- Tasks can be associated with methods of an enterprise bean using a task run-as policy. There are two run-as policies:

### **run as caller**

This is the default policy under which all methods run. If a request is already associated with a task, the configured method imports and runs under that task. If the request is not associated with a task, the request operates in the absence of a task and beans are controlled using the method-level access intent as configured outside of an application profile.

### **run as specified**

Under the run as specified policy, requests to the configured methods never run under an imported task; instead, a specified task names the unit of work beginning with the method.

- Tasks can be associated with any point in the execution of a component by using the programmatic interface of the task name manager. Any task that can be applied programatically must first be declared for the component; attempts to set unknown task names result in an `IllegalTaskNameException` error.

## **Application profiles**

An application profile is the set of access intent policies that should be selectively applied, as well as the list of tasks for which the policies should be applied.

The intention of application profiling is to enable an application to run under a different set of policies depending on the active task under which the application is operating.

Consider an application that centralizes the student records for a school district. These records are frequently accessed by the school district's central office in order to generate reports. The report generation process would be optimized if it held no locks with the back-end system, and if the records could be read into memory with as few back-end operations as possible. Occasionally, however, the records are updated by the students' instructors. Without the ability to distinguish between transactions, the developer is forced to assume a worst-case scenario and, wishing to use pessimistic concurrency, lock the records for all transactions.

Using the application profiling service, the developer can configure in as many ways as necessary the access intent under which the students' records are loaded. Under one profile, the records can be configured with an exclusive pessimistic update intent, not only locking-out competing transactions but ensuring that the student is not removed from the system before the transaction completes. Under another profile, the records can be configured with an optimistic intent as part of an object graph that is read from the back-end system in a single database operation. Any task configured with the pessimistic profile receives the strong-locking semantics required for certain transactions, while tasks configured with the optimistic profile receive the performance benefits appropriate for other transactions.

Multiple tasks can be configured on a single profile, indicating that different units of work can have the same requirements on the application; however, the same task cannot be registered with multiple application profiles because the run-time environment would have to guess which set of policies the developer wanted to have applied.

Use the Application Assembly Tool (AAT) to create and configure application profiles. Application profiles span the entire scope of an application. When a profile is created within a module, that profile is automatically created within all other EJB modules and at the EAR file level of the application. Likewise, when a profile is created at the EAR file level of an application, the profile is automatically pushed down to all EJB modules of the application. When an EJB module is imported into an application, all profiles within the module, and the profiles already declared inside the application, are merged.

## Application profiling performance considerations

Application profiling enables assembly configuration techniques that improve your application runtime, performance and scalability. You can configure tasks that identify incoming requests, identify access intents determining concurrency and other data access characteristics, and profiles that map the tasks to the access intents. The capability to configure the application server can improve performance, efficiency and scalability, while reducing development and maintenance costs. The application profiling service has no tuning parameters, other than a checkbox for disabling the service if the service is not necessary. However, the overhead for the application profile service is small and should not be disabled, or unpredictable results can occur.

Access intents enable you to specify data access characteristics. The WebSphere run-time environment uses these hints to optimize the access to the data, by setting the appropriate isolation level and concurrency. Various access intent hints can be grouped together in an access intent policy.

In WebSphere Application Server, access intent policies are configured on methods; when an invocation of a method drives the activation of a bean, the access intent configured on that method controls the data access characteristics, such as concurrency. In WebSphere Application Server Enterprise, application profiling enables you to configure multiple access intent policies on the entity bean. WebSphere Application Server Enterprise chooses the appropriate access intent for the entity bean, using the task associated with the current unit of work (UOW). Some callers can load a bean with the intent to read data, while others can load the bean for update. The capability to configure the application server can improve performance, efficiency, and scalability, while reducing development and maintenance costs.

Access intents enable the EJB container to be configured providing optimal performance based on the specific type of enterprise bean used. Various access intent hints can be specified declaratively at deployment time to indicate to WebSphere resources, such as the container and persistence manager, to provide the appropriate access intent services for every EJB request.

The application profiling service improves overall entity bean performance and throughput by fine tuning the runtime behavior. The application profiling service enables EJB optimizations to be customized for multiple user access patterns without resorting to "worst case" choices, such as pessimistic update on a bean accessed with the `findByPrimaryKey` method, regardless of whether the client needs it for read or for an update.

Application profiling provides the capability to define the following hierarchy: **Container Tasks > Application Profiles > Access Intent Policies > Access Intent Overrides**. Container tasks identify UOW and are associated with a method or a set of methods. When a method associated with the task is invoked, the task name is propagated with the request. For example, a UOW refers to a unique path within the application that can correspond to a transaction or activity session. The name of the task is assigned declaratively to a J2EE client or servlet, or to the method of an enterprise bean. The task name identifies the starting point of a call graph or subgraph; the task name flows from the starting point of the graph downstream on all subsequent IIOP requests, identifying each subsequent invocation along the graph as belonging to the configured task. As a best practice, wherever a UOW starts, for example, a transaction or an `ActivitySession`, assign a task to that starting point and make sure the task remains associated with that

UOW. It is best to configure the application server so that no more than one transaction is associated with more than one task.

The application profile service associates the propagated tasks with access intent policies. When a bean is loaded and data is retrieved, the characteristics used for the retrieval of the data are dictated by the application profile. The application profile configures the access intent policy and the overrides that should be used to access data for a specific task.

Access intent policies determine how beans are loaded for specific tasks and how data is accessed during the transaction. The access intent policy is a named group of access intent hints. The hints can be used, depending on the characteristics of the database and resource manager. Various access intent hints applied to the data access operation govern data integrity. The general rule is, the more data integrity, the more overhead. More overhead causes lower throughput and the opportunity for simultaneous data access from multiple clients.

If specified, access intent overrides provide further configuration for the access intent policy.

### **Best practices**

Application profiling is effective in a variety of different scenarios. The following are example situations where application profiling is useful

- **The same bean is loaded with different data access patterns**

The same bean or set of beans can be reused across applications, but each of those applications has differing requirements for the bean or for beans within the invocation graph. One application can require that beans be loaded for update, while another application requires beans be loaded for read only. Application profiling enables deploy time configuration for beans to distinguish between EJB loading requirements.

- **Different clients have different data access requirements**

The same bean or set of beans can be used for different types of client requests. When those clients have different requirements for the bean, or for beans within the invocation graph, application profiling can be used to tailor the bean loading characteristics to the requirements of the client. One client can require beans be loaded for update, while another client requires beans be loaded for read only. Application profiling enables deploy time configuration for beans to distinguish between EJB loading requirements.

### **Monitoring tools**

You can use the Tivoli Performance Viewer, database and logs as monitoring tools.

You can use the Tivoli Performance Viewer to monitor various metrics associated with beans in an application profiling configuration. The following sections describe at a high level the Tivoli Performance Viewer metrics that reflect changes when access intents and application profiling are used:

- **Collection scope**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor this information to determine the difference between using the ActivitySession scope versus the transaction scope. For the transaction scope, depending on how the container transactions are defined, activates and passivates can be associated

with method invocations. The application could use the `ActivitySession` scope to reduce the frequency of `activates` and `passivates`. For more information, see "Using the `ActivitySession` service."

- **Collection increment**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor *Num Activates* to watch the number of enterprise beans activated for a particular `findByPrimaryKey` operation. For example, if the collection increment is set to 10, rather than the default 25, the *Num Activates* value shows 25 for the initial `findByPrimaryKey`, before any result set iterator runs. If the number of `activates` rarely exceeds the collection increment, consider reducing the collection increment setting.

- **Resource manager prefetch increment**

The resource manager prefetch increment is a hint acted upon by the database engine to depend upon the database. The Tivoli Performance Viewer does not have a metric available to show the effect of the resource manager prefetch increment setting.

- **Read ahead hint**

The enterprise beans group contains EJB life cycle information, either a cumulative value for a group of beans, or for specific beans. You can monitor *Num Activates* to watch the number of enterprise beans activated for a particular request. If a read ahead association is not in use, the *Num Activates* value shows a lower initial number. If a read ahead association is in use, the *Num Activates* value represents the number of `activates` for the entire call graph.

**Database tools** are helpful in monitoring the different bean loading characteristics that introduce contention and concurrency issues. These issues can be solved by application profiling, or can be made worse by the misapplication of access intent policies.

Database tools are useful for monitoring locking and contention characteristics, such as locks, deadlocks and connections open. For example, for locks the DB2 Snapshot Monitor can show statistics for lock waits, lock time-outs and lock escalations. If excessive lock waits and time-outs are occurring, application profiling can define specific client tasks that require a more string level of locking, and other client tasks that do not require locking. Or, a different access intent policy with less restrictive locking could be applied. After applying this configuration change, the snapshot monitor shows less locking behavior. Refer to information about the database you are using on how to monitor for locking and contention.

The **application server logs** can be monitored for information about rollbacks, deadlocks, and other data access or transaction characteristics that can degrade performance or cause the application to fail.

---

## Assembling applications for application profiling

### Before you begin

Application profiling enables multiple access intent policies to be configured on the same entity bean, to be applied for a particular unit of work. Before using application profiles, you need to first apply access intent policies to entity beans. You can use the one of the default policies or create your own, as described in the topic, [Creating a custom access intent policy](#).

### Steps for this task

1. Configuring tasks.

Declaratively configure tasks using the Application Assembly Tool, as described in the following topics:

- Configuring a component task policy.
- Configuring a container task policy.

On rare occasions, you might find it necessary to configure tasks *programmatically*. Application profiling supports this requirement with a simple interface that enables both overriding of the current task associated with the thread of execution, and resetting of the current task to the original task. See the topic Using the TaskNameManager interface.

2. Creating an application profile.

3. Configuring tasks on application profiles.

## Applying access intent policies to entity beans

Use the Application Assembly Tool to apply an access intent policy to an entity bean, or set of entity beans, for a particular application profile. To configure access intent for beans that are to be loaded *without* application profiling support, apply access intent policies to your entity bean.

### Steps for this task

1. Start the Application Assembly Tool.

2. Create or edit the application EAR file.

For example, to change attributes of an existing application, select **File > Open**, then select the EAR file.

3. Select **EJB Modules > *module\_name* > Access Intent Policies**.

4. To configure a new access intent policy, right-click and select **New**.

5. On the **New Access Intent** panel, specify a name and description.

These attributes are provided as a convenience to the developer and is not used at run time.

6. To select the entities to which the access intent policy should apply, click **Add** beside the **Entity Beans** table.

7. From **Applied Access Intent** list, select an access intent policy.

8. To override an attribute defined in the applied policy, click **Add** beside the **Access Intent Attribute Overrides** table.

9. Click **OK** to exit the New Access Intent panel.

10. **(Optional)** Verify your configuration by selecting **File > Verify**.

11. Save your configuration by selecting **File > Save**.

### **Access intent assembly settings for application profiling**

Use this page to associate access intent policies with entity beans within an application profile in order to influence the management of persistence and collections.

**Name:** The name of the access intent policy.

An optional field provided for the convenience of the developer and not used by the run-time environment.

Data type

String

**Description:** A description of the access intent policy.

An optional field provided for the convenience of the developer.

Data type String

**Access intent entity beans:** The entity beans upon which to apply the access intent policy.

This policy is used to load the bean whether the bean is loaded by means of a finder, a container-managed relationship, or a dynamic query.

To add a new entity bean, click **Add** and select one or more entries from the list of eligible entities. To delete an entity, select it and click **Remove**. Entities will only be displayed for configuration if they have not already been configured within the same application profile. Also, only entities configured for version 2.x persistence are eligible for configuration.

**Applied access intent:** The access intent policy to apply to the selected entities.

From the drop-down menu, select the policy to apply to the selected entities. Predefined policies are available in addition to any custom policies that you have defined.

Default wsPessimisticUpdate - WeakestLockAtLoad



Range

**wsPessimisticUpdate - WeakestLockAtLoad**

Access type = Pessimistic update  
WeakestLockAtLoad = True  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsPessimisticUpdate**

Access type = Pessimistic update  
Collection scope = Transaction  
Collection increment = 1  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsOptimisticUpdate**

Access type = Optimistic update  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsPessimisticUpdate-Exclusive**

Access type = Pessimistic update  
Exclusive = true  
Collection scope = Transaction  
Collection increment = 1  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsPessimisticUpdate-NoCollision**

Access type = Pessimistic update  
No collision = true  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsPessimisticRead**

Access type = Pessimistic read  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsOptimisticRead**

Access type = Optimistic read  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**Custom access intent policies**

You can define any number of custom access intent policies.

**Access intent attribute overrides:** The values to override from the defined access intent policy.

The values configured in the defined access intent policy are shared by all referencing instances of applied access intents. Those values can be overridden here without affecting other configurations using the original intent.

To override an attribute, click **Edit**. To remove an existing attribute, select the attribute and click **Remove**.

**Collection scope:** If the EJB container returns a lazy collection, this intent specifies the maximum lifespan of that collection.

|         |  |
|---------|--|
| Default | Transaction  |
| Range   | <p><b>Transaction</b><br/>A lazy collection is valid for the life of the current transaction.</p> <p><b>ActivitySession</b><br/>A lazy collection is valid for the life of the current activity session. If collection scope is ActivitySession, but there is no activity session current at run time, the collection scope is automatically demoted to Transaction.</p> |

**Collection increment:** Specifies the number of objects to return in each segment of a lazy multi-object finder.

Lazy multi-object finders return  $n$  objects initially, and then repeatedly calls the server to fetch the next  $n$  objects as the client traverses the returned collection; this continues until the client has finished or the resultset is exhausted.

The value of  $n$  influences application performance; the collection increment count suggests the value of  $n$ . A value of 0 means return all elements (eager collection).

|           |   |
|-----------|---|
| Data type | Integer   |
| Default   | <p>Pessimistic update = 1<br/>Pessimistic update (exclusive = true) = 1<br/>All other access types = 25</p> |
| Range     | 0 - Integer.MAX_VALUE   |

**Resource manager prefetch increment:** Specifies the number of rows that the database is asked to return in each segment of a multi-object finder query.

This value is a hint to the JDBC driver; the actual behavior depends upon the specific driver implementation. A value of 0 indicates that the JDBC driver should ignore this attribute and make its own determination as to an appropriate fetch size.

|           |                       |
|-----------|-----------------------|
| Data type | Integer               |
| Units     | N/A                   |
| Default   | 0                     |
| Range     | 0 - Integer.MAX_VALUE |

**Read ahead hint:** For entities that have container-managed relationships (CMR), this field specifies which related objects should also be read in a single query when the `findByPrimaryKey()` method is executed.

The read ahead hint is given as a series of access paths composed of CMR relationship field names. Read ahead hints can only be specified if a single bean is selected for the policy and if the selected bean is configured with optimistic concurrency. CMR relationships are considered eligible for inclusion in a read

ahead hint only if the relationship is not many to many. Developers must ensure that all beans in a read ahead hint are loaded with optimistic concurrency at runtime; otherwise, an exception is raised when the CMR getter is invoked.

Data type

String

## Creating a custom access intent policy

Use the Application Assembly Tool to define a custom access intent policy, which can then be applied to an application's entity beans.

### Steps for this task

1. Start the Application Assembly Tool.
2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, select **File > Open**, then select the EAR file.
3. Select **EJB Modules > *module\_name* > Access Intent Policies**.
4. To create a new access intent policy, right-click and select **New**.
5. On the **New Defined Access Intent** panel, specify a name and a description.  
The name is a required attribute. It must be unique within the scope of the EJB module. The description is provided as a convenience to the developer and is not used at run time.
6. To define an access intent attribute, click **Edit** beside the **Access intent attributes** table.
7. From the pop-up menu, select an attribute.
8. Provide supported values in the new attribute panel, then click **OK**.
9. Click **OK** to exit the New Defined Access Intent panel.
10. **(Optional)** Verify your configuration by selecting **File > Verify**.
11. Save your configuration by selecting **File > Save**.

### **Method level access intent assembly settings**

Use this page to associate access intent policies with entity bean methods in order to influence the management of persistence and collections.

**Name:** The name of the access intent policy.

An optional field provided for the convenience of the developer and administrator.

Data type

String

**Description:** A description of the access intent policy.

An optional field provided for the convenience of the developer and administrator.

Data type

String

**Access intent methods:** The methods upon which to apply the access intent policy.

If the specified method causes the bean's data to be read from the backend, the configured access intent is associated with the bean instance for the duration of the unit of work; access intent policies configured on other methods are ignored. If the

invocation of the method results in the traversal of a CMR to a related entity, the configured access intent policy is used to load that related entity.

To add a new method, click **Add**. Expand the tree to select the method or methods from the EJB module. Be sure that each method has been configured with a policy no more than once. To delete a method, select it and click **Remove**. Note that only methods from 2.0 entity beans are available for configuration. If a CMP entity's version has been set to 1.1, the methods from that entity will not be visible in the method tree.

**Applied access intent:** The access intent policy to apply to the selected methods.

From the drop-down menu, select the policy to apply to the selected methods. Predefined policies are available in addition to any custom policies that you have defined.

Default          wsPessimisticUpdate - WeakestLockAtLoad

Range

**wsPessimisticUpdate - WeakestLockAtLoad**

Access type = Pessimistic update  
WeakestLockAtLoad = True  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsPessimisticUpdate**

Access type = Pessimistic update  
Collection scope = Transaction  
Collection increment = 1  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsOptimisticUpdate**

Access type = Optimistic update  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsPessimisticUpdate-Exclusive**

Access type = Pessimistic update  
Exclusive = true  
Collection scope = Transaction  
Collection increment = 1  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsPessimisticUpdate-NoCollision**

Access type = Pessimistic update  
No collision = true  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsPessimisticRead**

Access type = Pessimistic read  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**wsOptimisticRead**

Access type = Optimistic read  
Collection scope = Transaction  
Collection increment = 25  
Resource manager prefetch increment = 0  
Read ahead hint = null

**Custom access intent policies**

You can define any number of custom access intent policies.

**Access intent attribute overrides:** The values to override from the defined access intent policy.

The values configured in the defined access intent policy are shared by all referencing instances of applied access intents. Those values can be overridden here without affecting other configurations using the original intent.

To override an attribute, click **Edit**. To remove an existing attribute, select the attribute and click **Remove**.

**Collection scope:** If the EJB container returns a lazy collection, this intent specifies the maximum lifespan of that collection.

|         |  |
|---------|--|
| Default | Transaction  |
| Range   | <p><b>Transaction</b><br/>A lazy collection is valid for the life of the current transaction.</p> <p><b>ActivitySession</b><br/>A lazy collection is valid for the life of the current activity session. If collection scope is ActivitySession, but there is no activity session current at run time, the collection scope is automatically demoted to Transaction.</p> |

**Collection increment:** Specifies the number of objects to return in each segment of a lazy multi-object finder.

Lazy multi-object finders return  $n$  objects initially, and then repeatedly calls the server to fetch the next  $n$  objects as the client traverses the returned collection; this continues until the client has finished or the resultset is exhausted.

The value of  $n$  influences application performance; the collection increment count suggests the value of  $n$ . A value of 0 means return all elements (eager collection).

|           |  |
|-----------|--|
| Data type | Integer  |
| Default   | Pessimistic update = 1<br>Pessimistic update (exclusive = true) = 1<br>All other access types = 25 |
| Range     | 0 - Integer.MAX_VALUE  |

**Resource manager prefetch increment:** Specifies the number of rows that the database is asked to return in each segment of a multi-object finder query.

This value is a hint to the JDBC driver; the actual behavior depends upon the specific driver implementation. A value of 0 indicates that the JDBC driver should ignore this attribute and make its own determination as to an appropriate fetch size.

|           |                       |
|-----------|-----------------------|
| Data type | Integer               |
| Units     | N/A                   |
| Default   | 0                     |
| Range     | 0 - Integer.MAX_VALUE |

**Read ahead hint:** For entities that have container-managed relationships (CMR), this field specifies which related objects should also be read in a single query when the `findByPrimaryKey()` method is executed.

The read ahead hint is given as a series of access paths composed of CMR relationship field names. Read ahead hints can only be specified if a single bean is selected for the policy and if the selected bean is configured with optimistic concurrency. CMR relationships are considered eligible for inclusion in a read

ahead hint only if the relationship is not many to many. Developers must ensure that all beans in a read ahead hint are loaded with optimistic concurrency at runtime; otherwise, an exception is raised when the CMR getter is invoked.

Data type String

### Defined access intent assembly settings for EJB modules

Use this page to create custom access intent policies.

Access intent policies are created by developers to be applied to beans across the EJB module. Changes made to a defined policy upon which beans have previously been configured are reflected in the behavior of those configured beans.

**Name:** The name of the access intent policy used to identify the attributes to be applied to methods.

The name must be unique within the EJB module.

Data type String

**Description:** A description of the access intent policy.

An optional field provided for the convenience of the developer.

Data type String

**Access intent attributes:** To modify an attribute, click **Edit**.

**Access type:** Specifies whether the concurrency control scheme used to select and update the bean in the current transaction should be pessimistic or optimistic.

The same scheme must be used in a given transaction. Any attempt by the application to suggest a change in scheme within a given transaction is ignored. Also specifies whether the application intends to read or update the entity bean. This hint is used to determine isolation level on JDBC backend connections, as well as lock type shared or exclusively used by the EJB container to manage access to enterprise beans with an activation policy of once.

Default Pessimistic update

## Range

### Optimistic read, Pessimistic read

A value of read indicates that the caller does not intend to drive an update method on the entity bean. If a read intent is indicated and the caller drives a method that performs an update, the following result occurs

- For a CMP entity, the EJB Container throws `com.ibm.websphere.pm.UpdateMethodCannotProceedWithIntegrity`;
- for a BMP entity, the EJB is advised to throw `javax.ejb.EJBException`

### Optimistic update, Pessimistic update

A value of update indicates that the caller intends to drive update methods on the entity bean. In addition, you can indicate *one* of the following hints to further qualify the pessimistic update attribute:

#### Weakest lock at load

The application prefers the weakest lock available from the back-end datastore, assuming that the primary access of data will be for read purposes; updates can result in lock escalation if necessary.

**Data type**  
Boolean

**Units** N/A

**Default**  
Selected (True)

**Range** True, False

#### Exclusive

The application prefers exclusive access to the database rows. A higher isolation should be chosen by the run-time environment.

**Data type**  
Boolean

**Units** N/A

**Default**  
Cleared (False)

**Range** True, False

#### No collision

The application has no row collisions by design; the WebSphere run-time environment takes no responsibility for ensuring this condition, and can choose a lesser isolation level.

**Data type**  
Boolean

**Default**  
Cleared (False)

**Range** True, False



**Collection scope:** Specifies the maximum lifespan of a lazy collection when such a collection is returned by the EJB container

Default  
Range

Transaction

**Transaction**

A collection scope of Transaction specifies that a lazy collection is valid for the life of the current transaction.

**ActivitySession**

A collection scope of ActivitySession specifies that a lazy collection is valid for the life of the current activity session. If collection scope is ActivitySession, but there is no activity session current at runtime, then the collection scope is automatically demoted to Transaction.

**Collection increment:** Specifies the number of objects to return in each segment of a lazy multi-object finder.

Lazy multi-object finders return  $n$  objects initially, then repeatedly return to the server to fetch the next  $n$  objects as the client traverses the returned collection; this continues until the client is finished or the resultset is exhausted. The value of  $n$  influences application performance; the collection increment count suggests the value of  $n$ . A value of 0 means return all elements (eager collection).

Data type  
Default

Integer

Pessimistic update = 1  
Pessimistic update (exclusive = true) = 1  
All other access types = 25

Range

0 - Integer.MAX\_VALUE

**Resource manager prefetch increment:** Specifies the number of rows that the database is asked to return in each segment of a multi-object finder query.

This value is a hint to the JDBC driver; the actual behavior depends upon the specific driver implementation. A value of 0 indicates that the JDBC driver should ignore this attribute and make its own determination as to an appropriate fetch size.

Data type  
Default  
Range

Integer

0

0 - Integer.MAX\_VALUE

## Configuring a component task policy

Use the Application Assembly Tool to configure a component's own task. A servlet's or application client's own task is associated with distributed requests from the component unless overridden by a container task policy. Components without a configured task of their own run without a task; units of work without a task and application profile load entities using the method-level or default access intent configuration, without the support of application profiling.

### Steps for this task

1. Start the Application Assembly Tool.

2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, select **File > Open** then select the EAR file.
3. Select **Web Modules > *module\_name.war* > Web Components > *component\_name***
4. Select the **WAS Enterprise** tab.
5. Select the **Own task** checkbox. Provide a name and description for the task.  
Task names do not have to be unique within an application; however, task names should be shared consciously and conservatively; at run time, all tasks with the same name are treated the same way, regardless of where the task was configured.  
The description is provided as a convenience to the developer and is not used by the run-time environment.
6. Click **OK**.
7. **(Optional)** Select **File > Verify** to verify your configuration.
8. Select **File > Save** to save your configuration.

#### What to do next

Configuring tasks on application profiles.

## Configuring a container task policy

Use the Application Assembly Tool to apply a container task policy to a method, or set of methods, for an application's entity beans. A container task policy defines the task under which the method is invoked. A method can run with an imported task, with its own task, or with a specified task.

**Note:** Applications use application profiling on a *per task* basis. As soon as a path within an application is configured with a task and that task is associated with an application profile, any entity enlisted within units of work in the scope of that task cease to use the method-level access intent configuration and instead use the bean-level access intent configuration of the application profile with which the task is registered. If the entity was not configured with an intent within the application profile, then the configured default intent is applied.

#### Steps for this task

1. Start the Application Assembly Tool.
2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, select **File > Open** then select the EAR file.
3. Select **EJB Modules > *module\_name* > Container Tasks**.
4. To create a new container task policy, select **File > New > Selected Object**.
5. On the **New Container Task** panel, specify a name and a description.  
These attributes are provided as a convenience to the developer and are not used at runtime.
6. To select the methods to which the container task policy should apply, click **Add** beside the **Methods** table.
7. Select one of the these attributes:

#### **Run as caller**

This is the default attribute applied to EJB method invocations. If the configured methods are invoked with an associated task, the method

is executed with the imported task. If a method is invoked by a request that is unassociated with a task, the method continues to execute without a task.

#### **Run as specified**

The configured methods are never invoked with an imported task. Instead, the method executes as the specified task name. The task name can be selected from the pull-down menu or you can enter a new task name. The description is provided as a convenience to the developer and is not used at runtime.

8. Click **OK**.
9. **(Optional)** Select **File > Verify** to verify your configuration.
10. Select **File > Save** to save your configuration.

#### What to do next

Configuring tasks on application profiles.

### **Container assembly settings for tasks**

Use this page to configure container task policies.

Tasks identify units of work within a distributed application by the implicit propagation of the task name on remote requests. The task name is configured on application profiles in order to customize the access intent for the units of work associated with that task.

A container task policy instructs the container under which task a request upon an EJB method should operate. Methods can be configured to run as a caller's task, as the task configured on the bean, or as a specified task.

**Name:** The name of the policy.

An optional field provided for the convenience of the developer. Although Name is a required field, the name is not used except as a label within the application assembly tool.

Data type String

**Description:** A description of the policy.

An optional field provided for the convenience of the developer.

Data type String

**Methods:** The methods upon which the container will apply the task policy.

To add a new method to the policy, click **New**. Expand the tree to select the method or methods from the EJB module. Be sure that each method has been configured no more than once with a container task policy. To remove a method, select it and click **Remove**.

**Name:** The policy that the container should apply when the configured set of methods are invoked.

Default Run as caller

Range

**Run as caller**

If the client invokes the bean method with an associated task, the container invokes the bean method with the same task. If the client invokes the bean method without a task, the container invokes the bean method with the task configured as the bean's default task.

**Run as specified**

The container invokes the bean method with the specified task.

**Name** The name of the task

**Description**

The description of the task

## Creating an application profile

Use the Application Assembly Tool to create an application profile. An application profile contains a set of access intent policies applied to an application's entity beans. The access intent policies are only applied for requests that are associated with tasks configured on the application profile.

### Steps for this task

1. Start the Application Assembly Tool.
2. Create or edit the application EAR file.

For example, to change attributes of an existing application, select **File > Open**, then select the EAR file.

3. You can create application profiles at the scope of either an application EAR file or an EJB module.

To create an application profile at the EAR file scope, select **Application Profile**.

To create an application profile at the module scope, expand **EJB Modules > module\_name > Application Profile**.

An application profile logically spans the application EAR file and all contained EJB Jar files. When a profile is created for the EAR file, the profile is automatically created within all EJB Jar files. When a profile is created within an EJB Jar file, the profile is automatically created for all remaining EJB Jar files and the EAR file as well. It makes no difference where an application profile is created, edited, or removed.

4. To create a new application profile, select **File > New > Selected Object**.
5. On the **New Application Profile** panel, specify a name and a description.

The name of the profile must be unique within the application; there cannot be two distinct application profiles with the same name.

The description is provided as a convenience to the developer and is not used at runtime.

6. Click **OK**.
7. Be sure that your newly-created application profile is selected. In the navigation pane, expand the new application profile, then select the **Access Intent** node.

8. To apply access intents within the scope of the application profile, follow the steps as described in the topic Applying access intent policies to entity beans. Any custom access intent policies are available within the application profile.
9. Select **File > Verify** to verify your configuration.
10. Select **File > Save** to save your configuration.

## Application profile assembly settings

Use this page to configure application profiles.

Application profiles support the definition of alternate access-intent configurations that are mapped to particular requests identified by an association with a task name.

**Name:** The name of this application profile.

The name must be unique; multiple profiles cannot share the same name.

The creation, configuration, and deletion of profiles is reflected in the configuration of profiles at both the application and module scope

Data type String

**Description:** A description of the application profile

An optional field provided for the convenience of the developer and administrator.

Data type String

**Tasks:** Tasks that are configured to operate under the application profile.

Requests associated with any of the configured tasks operate under the access-intent policies configured with the profile within the EJB modules.

To add a task that has been declared within the application, click **Add**. To add a task that has not been declared within the application, click **New**.

To remove a task, select the task and click **Remove**. Any given task can be configured on only *one* application profile.

**Name:** The name of the task.

Select the name of the task from the pulldown menu or specify a new name. The name of the task must be unique among the set of application profiles. The task name is a required field.

Data type String

**Description:** A description of the task.

An optional field provided for the convenience of the developer and administrator.

Data type String

## Configuring tasks on application profiles

Use the Application Assembly tool to associate tasks with application profiles. When a task is configured on an application profile, the access intent policies defined in the profile are applied, as appropriate, to requests associated with that task.

### Steps for this task

1. Start the Application Assembly Tool.
2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, select **File > Open**, then select the EAR file.
3. Associate tasks with an application profile at the scope of either an application EAR file or an EJB module.  
To edit the tasks within an application profile at the EAR file scope, select **Application Profile** in the navigation pane, then select the appropriate profile from the list of application profiles.  
To edit the tasks within an application profile at the EJB module scope, expand **EJB Modules > module\_name > Application Profile**, then select the appropriate application profile from the list of application profiles.  
An application profile logically spans the application EAR file and all contained EJB Jar files. Adding a task to, or removing a task from, an application profile within any module or EAR file automatically causes the task to be removed from the same profile in other modules.
4. To add a task defined elsewhere in the application, click **Add** beside the table of tasks and select a task name from the dropdown menu. To add a task that you defined outside of the application, click **New** beside the table of tasks.
5. Select **File > Verify** to verify your configuration.
6. Select **File > Save** to save your configuration.

## Dynamic query assembly settings

Use this page to associate access intent policies with entity beans in order to influence the management of persistence during the execution of dynamic queries. This access intent will be used for requests that are not associated with an application profile.

### **Name**

The name of the access intent policy. The name is provided as a convenience to the developer and is not used by the application server.

Data type String

### **Description**

A description of the access intent policy. An optional field provided for the convenience of the developer.

Data type String

### **Entity beans**

The beans with which to associate the access intent.

To add a new bean, click **New**. Expand the tree to select the bean or beans from the EJB module. Be sure that each bean has been configured with a policy no more

than once within any particular application profile. To delete a bean, select it and click **Remove**. Note that only 2.0 entity beans are available for configuration. If a CMP entity bean's version has been set to 1.1, that bean is not visible in the tree.

### Applied access intent

The access intent policy to apply to the selected beans.

Select from the pulldown menu the policy to apply to the selected beans. The following policies are available:

|         |  |
|---------|--|
| Default | wsPessimisticUpdate - WeakestLockAtLoad  |
| Range   | <b>wsPessimisticUpdate - WeakestLockAtLoad</b><br>Access type = Pessimistic update<br>WeakestLockAtLoad = True |
|         | <b>wsPessimisticUpdate</b><br>Access type = Pessimistic update   |
|         | <b>wsOptimisticUpdate</b><br>Access type = Optimistic update   |
|         | <b>wsPessimisticUpdate-Exclusive</b><br>Access type = Pessimistic update<br>Exclusive = true                   |
|         | <b>wsPessimisticUpdate-NoCollision</b><br>Access type = Pessimistic update<br>No collision = true              |
|         | <b>wsPessimisticRead</b><br>Access type = Pessimistic read   |
|         | <b>wsOptimisticRead</b><br>Access type = Optimistic read   |

---

## Managing application profiles

Manage your application profiles using the administrative console. From the console, you can add tasks to, and remove tasks from, application profiles.

### Steps for this task

1. Start the administrative console.
2. Select **Applications > Applications > *application\_name* > Application Profile > *profile\_name* > Tasks**.
3. On the Tasks collection page, you can add new tasks to the profile, delete tasks, edit current task settings, and so on.

Note that no task can, within the scope of an application, be configured on more than one application profile. In such a situation, your application cannot be restarted until you correct the configuration.

4. Save your configuration.
5. Restart the application in order for your changes to take affect.

## Application profiling exceptions

The following exceptions are thrown in response to various illegal actions related to application profiling:

### **com.ibm.ws.exception.RuntimeWarning**

This exception is thrown when the application is started, if the application is configured incorrectly. The startup is consequently terminated. You can validate an application's configuration by using the **Verify** function in the Application Assembly Tool. Some examples of misconfiguration include:

- A task configured on two different application profiles.
- A method configured with two different task run-as policies .

### **com.ibm.websphere.appprofile.IllegalTaskNameException**

This exception is raised if an application attempts to programmatically set a task when that task has not been configured as a task name reference.

## Application profiling service settings

Use this page to enable or disable the application profiling service.

Applications that are configured to use the application profiling service will not start successfully unless the application profiling service is enabled.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Application Profiling Service**.

### **Startup**

Specifies whether the server will attempt to start the application profiling service.

|         |  |
|---------|--|
| Default | Selected   |
| Range   | <b>Selected</b><br>When the application server starts, it attempts to start the application profiling service automatically. |
|         | <b>Cleared</b><br>This option is unavailable. The application profiling service cannot be disabled.                          |

## Application profile collection

Use this page to manage application profiles.

An application profile is a set of policies that are to be applied during the execution of an enterprise bean and a set of tasks that are associated with that profile. Mapping tasks to application profiles will control which access intent policies are applied at run time for the units of work that correspond to a particular task.

To view this administrative console page, click **Applications > Applications > *application\_name* > Application Profile**.

### **Name**

The name of the application profile.

The name must be unique; multiple profiles cannot share the same name.



Data type String

### Description

A description of the application profile.

Data type String

### Application profile settings

Use this page to modify application profile settings.

To view this administrative console page, click **Applications > Applications > *application\_name* > Application Profile > *application\_profile\_name***.

**Name:** The name of the application profile.

The name must be unique; multiple profiles cannot share the same name.

Data type String

**Description:** A description of the application profile.

Data type String

### Task collection

Use this page to manage tasks.

Requests associated with any of the configured tasks operate under the access-intent policies that are configured with the profile. A task can be configured on only *one* application profile.

To view this administrative console page, click **Applications > Applications > *application\_name* > Application Profile > *application\_profile\_name* > Tasks**.

**Name:** The name of the task.

The task name must be unique among the set of application profiles.

Data type String

**Description:** A description of the task.

Data type String

### Task settings

Use this page to modify task settings.

To view this administrative console page, click **Applications > Applications > *application\_name* > Application Profile > *application\_profile\_name* > Tasks > *task\_name***.

**Name:** The name of the task.

The task name must be unique among the set of application profiles.

Data type String

**Description:** A description of the task.

Data type String

---

## Using the TaskNameManager interface

You can declaratively configure tasks on a J2EE component and associate tasks with EJB methods using the Application Assembly Tool. On rare occasions, you might find it necessary to *programmatically* set the current task name. Application profiling supports this requirement with a simple interface that enables both overriding of the current task associated with the thread of execution, and resetting of the current task with the original task.

Application profiling does not support queries of the task that is in operation at run time. Instead, applications interact with logical task names that are declaratively configured as task references. Logical references enable the actual task name to be changed without having to recompile applications.

While you cannot programmatically set the current access intent policy, you can accomplish this by programmatically setting a task. Consider, for example, an entity bean with a single multi-object finder method, `getLargeAccounts()`, which is invoked by the `single()` method of the `AccountManager` session bean. By default, suppose that the entity bean runs assuming read intent; suppose, also, that an application profile has been configured under which that bean loads assuming update intent. Configure a task, perhaps with the name "update", on the profile. Now, depending on logic in the session bean's method, the session bean selectively and programmatically sets the task update before invoking the entity bean's finder method; that finder method correctly functions assuming updates.

Wherever possible, avoid setting tasks programmatically. The declarative method results in more portable function that can be easily adjusted without requiring redevelopment and recompilation.

### Steps for this task

1. Configure task references.

Application profiling requires that a task name reference be declared for any task that is to be set programmatically. Task name references introduce a level of indirection so that the actual task set at run time can be adjusted by reassembly without requiring recoding or recompilation. Any attempt to set a task name that is undeclared as a task reference results in the raising of an exception.

- a. Start the Application Assembly Tool.
- b. Create or edit the application EAR file.  
For example, to change attributes of an existing application, select **File > Open**, then select the EAR file.
- c. Select **Web Modules > module\_name.war > Web Components > component\_name**
- d. Select the **WAS Enterprise** tab.
- e. Click **Add** beside the **Task references** table.
- f. Provide a name for the task reference.

The name must be unique for the component. This is the name that is referenced programatically. The name should be short and should be descriptive of the function that is performed when the task is executed.

- g. Provide a name and description for the task, itself.

The name of the task is the identifier that is propagated on remote requests; it is the task name that is configured on application profiles to dynamically associate access intent hints with entity bean execution.

- h. Click **OK**.

- i. Select **File > Save** to save your configuration.

2. Perform a JNDI lookup on the TaskNameManager interface:

```
InitialContext ic = new InitialContext();
TaskNameManager tnManager = ic.lookup
("java:comp/websphere/AppProfile/TaskNameManager");
```

The TaskNameManager interface is not bound into the namespace if the application profiling service is disabled.

3. Set the task name:

```
try {
tnManager.setTaskName("updateAccount");
}
catch (IllegalTaskNameException e) {
// task name reference not configured. Handle error.
}
// . . .
rnManager.resetTaskName();
```

Resetting the task name undoes the effects of any setTaskName() method operations and reestablishes whatever task name was current when the component began execution. If the setTaskName() method has not been called, the resetTaskName() method has no effect.

### What to do next

Configuring tasks on application profiles.

## **TaskNameManager interface**

The TaskNameManager interface is available to all J2EE components using the following JNDI lookup:

```
java:comp/websphere/AppProfile/TaskNameManager
package com.ibm.websphere.appprofile;
```

```
/**
 * The TaskNameManager is the programmatic interface
 * to the application profiling function. Using this interface,
 * programmers can set the current task name on the
 * thread of execution. The task name must have been
 * configured in the deployment descriptors as a task
 * reference associated with a task. The set task
 * name's scope is the duration of the method
 * invocation in the EJB and Web components and for
 * the duration of the client process, or until the
 * resetTaskName() method is invoked.
 */
public interface TaskNameManager {

/**
 * Set the thread's current task name to the specified
 * parameter. The task name must have been configured as
```

```
* a task reference with a corresponding task or the
* IllegalArgumentException exception is thrown.
*/
public void setTaskName(String taskName) throws IllegalArgumentException;

/**
 * Sets the thread's task name to the value that was set
 * at, or imported into, the beginning of the method
 * invocation (for EJB and Web components) or process
 * (for J2EE clients).
 */
public void resetTaskName();
}
```

---

## Chapter 16. Using Business Rule Beans

### Before you begin

This topic provides a brief overview of the steps involved in externalizing Business Rule Beans. To gain an understanding of business rules and Business Rule Beans (BRBeans), review the topic Overview of Business Rule Beans. The following sections provide an overview of externalizing business rules using Business Rule Beans:

### Steps for this task

1. Developing BRBeans
2. Assembling applications for use with BRBeans
3. Managing rules

### Usage scenario

To help you get started quickly, a sample BRBeans application is provided. Samples are installed by default during a typical WebSphere installation or you can select to install specific samples during a custom installation.

The BRBeans sample is an online movie store application. The application's EAR file is installed and the application is configured to use the Cloudscape database (Cloudscape is provided with WebSphere Application Server). A number of rules are created that you can view using the Rule Management Application (RMA). To do this, change to the `<install_root>/bin` directory and type one of the following commands:

- On a Microsoft Windows platform:  

```
rulemgmt ..\samples\lib\BRBeans\movieSampleProperties
```
- On a Unix platform:  

```
rulemgmt.sh ../samples/lib/BRBeans/movieSampleProperties
```

By running the sample, you can see how these rules are used. The source code for the sample also is provided in the `<install_root>/samples/src/BRBeans/Movie` directory. To see the use of trigger points, search the code for places where the `trigger()` method is used.

---

## Advantages of externalizing business rules

Business Rule Beans (BRBeans) provide a framework in which business applications can externalize business rules. You can externalize rules by extending your application analysis and design processes to identify the points of variability (or "trigger points") in application behavior. When the application is implemented, the business logic required at the points of variability is externalized into a business rule. This allows certain aspects of the behavior to be changed without actually changing the application code.

Here are some advantages of externalizing business rules:

**Explicit documentation of business practice decisions**

Separating business rule values from the application code makes the code easier for others to view and understand while isolating information that relates to business practice decisions.

**Clearer understanding of application behavior**

Externalization makes it possible to inspect the application to see which business rules are being applied, when they are applied, and under what circumstances.

**Reuse of rules across business processes**

Separating rules from the business logic of the application makes it easy to reuse a business practice decision in a consistent fashion.

**Increased consistency of business practices**

Because externalized rules promote reuse and facilitate clear understanding of business practice decisions, they provide a basis for improving business practice consistency across applications.

**Decreased maintenance and testing costs**

Externalized rules have a clearly defined scope and are not tightly coupled to the application code. This makes them easy to modify, quick to test, and decreases costs and improves cycle time.

**Improved manageability of business practice decisions**

Externalization, change history, and inspectability all promote clear ownership and consequently a better definition of who can change rules and under what circumstances.

**Increased confidence in predicting the business impact of proposed changes**

Because rules are available for inspection, have well-defined scope, and are not tightly coupled to application business logic, they make it easy to understand the likely impact of changes and to predict whether contemplated modifications or additions will have unwanted ripple effects.

**Ability to identify and correct conflicting business rules in different parts of the business**

Externalized rules make it easy to check that rules being used in two different parts of an application or even two different applications dealing with different parts of the business, are consistent.

---

## Overview of Business Rule Beans

Business Rule Beans are used to create and modify rules that keep pace with complex business practices. This enables your application's core behavior and user interface objects to remain intact and untouched, even as business practices change.

The Business Rule Beans (BRBeans) framework enables you to organize rules in folders. Folders provide a structure similar to the file system on your computer's hard drive. For example:

- Rules can be placed in folders based on any criteria you want.
- A rule folder can contain any number of rules and other rule folders.

In the BRBeans framework, each business rule is represented by an entity bean that persistently stores information related to that rule. Each business rule is assigned an appropriate rule name and stored in an appropriate rule folder (See Rule folders for more information).

When naming rules and folders, adhere to the Java package naming convention. That is, name rules and folders based on the domain name of the organization for which the rules are developed. For example, ACME's `isSeniorCitizen` rule's fully qualified rule name ("full rule name"), might be `com/acme/ageRules/isSeniorCitizen`. In this example, the `com/acme` path is used by all of the rules developed by ACME and the `ageRules` folder is used to separate "age" rules from rules of other kinds. The root folder has no name; therefore, fully qualified path names never start with a forward slash (`/`).

A fully qualified rule name consists of the following:

- The full path of the folder followed by a forward slash (`/`)
- The name of the rule

This fully qualified rule name is used by a trigger point to identify the rule to trigger. Trigger points are small pieces of code that interface with the Business Rule Beans trigger point framework to run business rules during application execution. See [Placing a trigger point](#) for more information.

By default, trigger points can only trigger rules that are currently in effect based on the current date and time when the trigger point is called. A business rule has a start date and an end date (see [Rule attributes](#) for more information) that together define the interval during which the rule is in effect (see [Rule states](#) for more information). This behavior can be overridden by specifying a date on the trigger point. This date is referred to as the "As Of Date". If no start date is specified, the rule is not valid and cannot be found by trigger points. Conversely, if no end date is specified, the rule never expires. Dates and times with a precision of one second can be assigned using the "Rule Management Application".

When there is more than one rule with the same fully qualified name, all of the rules with that name that are currently in effect are triggered and the results are combined using the combining strategy specified on the trigger point. See the [CombiningStrategy](#) method for more information.

## Externalized business rules

A business rule is a statement that defines or constrains some aspect of a business by asserting control over some behavior of that business.

A business rule officiates over frequently changing business practices and can come from within the company or be mandated from outside, typically by regulatory agencies. Typical uses for business rules include the following:

- Determining the current interest rate for a home loan
- Calculating a discount for a product
- Calculating the tax to apply to a given product
- Determining whether a given person is a senior citizen

The objects used to implement a business rule contain methods and attributes used by the Business Rule Beans (BRBeans) run-time environment, its administrative component, or both. An externalized business rule is implemented as a pair of objects:

- Rule
- RuleImplementor

The **Rule** is an entity enterprise bean that stores all of the persistent data for the business rule. This is the object that the trigger point framework code actually deals with directly. When a trigger point is invoked, the internal framework code performs a query to find the Rule object or objects representing the business rules to be triggered. Once the Rules are found, the framework code determines where the Rule is invoked, either local to the trigger point or remotely on the application server. Then, it invokes the `fire` method on either the Rule enterprise bean itself (for remote triggering) or on a local copy of the enterprise bean (for local triggering) to perform the function of the business rule.

The class name of the business rule's **RuleImplementor** is stored persistently in the Rule. The `RuleImplementor` is a transient object (not managed by the application server) that the Rule instantiates and then uses to do the actual work. When the `fire()` method is called on the Rule object, the Rule object combines its persistent set of values with the parameters it received on invocation. This creates the parameter list for the `RuleImplementor` prior to invoking `fire()` on the `RuleImplementor` with this parameter list. The actual execution of the `RuleImplementor` algorithm can take place either remotely (within the application server where the BRBeans enterprise beans are installed) or locally (within the Java virtual machine (JVM) where the trigger point was called).

## Types of business rules

Business rules can be divided into the following types:

- Base rules
- Classifier rules

**Base rules** are the most common type of rule and are triggered by the `TriggerPoint.trigger` method. You can divide Base rules into the following categories:

### Derivation rules

These rules use an algorithm to return a value. These rules return any type of value that makes sense in the business context in which they are used. For example, a derivation rule can calculate a discount or compute the total price of an order.

### Constraint rules

These rules confirm that an operation has met all of its obligations and that a particular constraint or edit has been met. For instance, a constraint rule can check that a value entered by an external user is within legal bounds. Business Rule Beans (BRBeans) provide a special return type, `com.ibm.websphere.brb.ConstraintReturn`, which can be returned by a constraint-type rule. A `ConstraintReturn` object contains a boolean value so that if it is false, it can contain information that can be used to produce an external message explaining what constraint was not met.

### Invariant rules

These rules ensure that multiple changes made by an operation are properly related to one another.

### Script rules

These rules implement "micro-workflow" or electronic performance support. They are small, variable pieces of a business process that provide assistance to end-users to get the most from the application.



On the surface, **classifier rules** are much like base rules. However, classifier rules can be used to determine the ways in which variables are classified by a business. Classifier rules are triggered by the `TriggerPoint.triggerClassifier` method.

A classifier rule is used to compute a classification for a particular business situation. The classification returned is required to be of type `string`. For instance, bank customers may be classified into gold, silver, and bronze categories based on their spending history or the amount of money they have in their account. For more information on this type of rule, refer to Situational trigger point.

## Rule folders

Rule folders are similar to the directories that divide a computer's hard drive in that they split a large number of files into conceptual units. The rule folder adds its path to the fully qualified rule name. Like the directories on a hard drive, a rule folder can contain any number of rules or rule folders.

Although you can name the folders whatever you deem appropriate, it is recommended that you follow the Java package naming convention. That is, base the names on the domain name of the organization where the rules are developed. So, the fully qualified rule name or full rule name of ACME's `isSeniorCitizen` might be `com/acme/ageRules/isSeniorCitizen`. In this example, the `com/acme` path is used by all of the rules developed by ACME and the `ageRules` folder is used to separate "age" rules from rules of other kinds.

**Note:** The root folder has no name, which means that fully qualified path names never start with a `'/'`.

When using the Rule Management APIs, a rule folder contains instances of `IRules`, which also are referred to as "rules". To begin working with rules, get the root rule folder by using the `getRootFolder` method on `RuleMgmtHelper` class. From the root rule folder you can add, delete, and retrieve folders and rules using methods on this interface.

## Rule attributes

### Rule name

A name for the rule that is appropriate to its business context.

### Rule folder

The folder that contains the rule.

### Start date

This is the date and time at which the rule goes into effect. Prior to this time, it will not be found by trigger points. Together with the end date, the start date defines a period of time during which the rule is effective. A rule that does not have a start date specified is not a valid rule and will not be found by trigger points.

### End date

This is the date and time at which the rule is no longer effective. After this date and time the rule is no longer in effect and will not be found by trigger points. Together with the start date, the end date defines a period of time during which the rule is effective. A rule that does not have an end date specified is valid and will never expire.

### Ready

This indicates whether the rule is ready to be used. Rules that are not marked as ready will not be found by trigger points. This is intended to be

an easy way to keep a rule from being used until it is completely defined or to temporarily turn a rule off without having to change the basic rule data such as start and end dates.

#### **Java Rule Implementor name**

This is the fully package-qualified name of a Java class that implements the BRBeans RuleImplementor interface. The fire method of the class performs the function of the rule. Business Rule Beans (BRBeans) provide several predefined rule implementors or you can write your own. See Rule Implementors or Customized rule implementors for more information.

#### **Initialization parameters**

This is an array of parameters that are passed to the rule implementor to initialize it. Each element in the array can be any object. This also can be referred to as the rule data, which is the external data that may change over time. The initialization parameters defined for a rule are passed directly to the init method of the rule implementor when it is instantiated. See "Rule Implementors" for more information on how rule implementors can use initialization parameters.

#### **Firing parameters**

Normally, firing parameters are simply the parameters passed on the trigger point when a rule is triggered. However, it is allowed to override these parameters by specifying parameters on the rule itself. This is where these overriding parameters are specified.

#### **Firing location**

This specifies where the rule implementor for this rule is instantiated and run. The following values are allowed:

**Local** This option instantiates the rule implementor and runs it local to the trigger point (in the same JVM as the trigger point call). This is run on the client machine if the trigger point call is done there or on the server if the server part of an application makes a trigger point call. Use this option for the best performance since, once a rule is cached on the client, the entire triggering process can be performed locally without going to the server at all. The main disadvantage of this option is that the class files for the rule implementors need to be available on every client that can trigger rules.

#### **Remote**

This will instantiate the rule implementor and run it on the application server where the Business Rule Beans enterprise beans are installed. When using this option at least one remote method call always is required to trigger a rule since the trigger takes place on the server. The advantage is that the rule implementor class files only need to be available on the server.

#### **Anywhere**

This option tries to instantiate and run the rule implementor locally, and, if the class cannot be found, it tries to trigger it remotely.

#### **Classification**

For classified rules, this is the classification to which the rule applies. This is used when you use a situational trigger. Once a classification is

computed for the situational trigger point, rules that apply to that classification are found and triggered. For more information, see Situational trigger point.

**Classifier**

This indicates whether this rule computes a classification. Classification is used for a situational trigger. A classifier rule is used to perform the first step of a situational trigger which computes a classification that is used to find rules to deal with the situation. For more information, see "Situational trigger point".

**Dependent rules**

In many cases, a rule triggers other rules to complete the overall task. These other rules are referred to as dependent rules and can be specified using the dependent rules attribute. For more information, see Dependent rules.

**Business intent**

This is a text description of the intent of this rule from the view point of the business analyst. You can store any text string here.

**Description**

This is a text description of the rule at the programmer's level. You can store any text string here.

**Original requirement**

This is a text description of the initial business analyst requirement of this rule. You can use this description to keep track of why this rule was originally created (for example, to keep auditing records). You can store any text string here.

**User-defined data**

You can store a user-defined text string here. The format and use of this data is completely determined by the user.

**Primary key**

Every rule has a primary key to uniquely identify it in the database where the enterprise beans are stored. Normally, a unique primary key is generated automatically when you create a new rule. However, you can use the rule management APIs to specify your own primary key, if desired. See Rule management APIs for more information.

**Precedence**

This is the relative priority of this rule. The default finding strategy uses this value to order the rules found in the database, from lowest to highest, when more than one rule is found for a particular trigger point. Rules are sorted numerically by precedence with the numerically lowest precedence first and the numerically highest precedence last.

## Rule states

Rules can be in any one of the following states at any particular time:

**scheduled**

The rule is scheduled to become effective (its start date is in the future) and will not be found by current trigger points.

**in effect**

The rule is currently in effect and can be found by trigger points.

**expired**

The rule is no longer in effect (the end date is in the past) and will not be found by trigger points.

**invalid**

The rule is not correctly defined and will not be found by trigger points.

Typically, only those rules that are "in effect" are found by the Business Rule Beans (BRBeans) run-time environment. This behavior can be overridden by setting an `asOfDate` on the `TriggerPoint` object, which then will execute "as if" the current date is the given date. For more information, see "As Of Date".

When a Rule is first created, it is marked as "ready for use" and is found when firing Rules. If the Rule is not complete and you do not want it to be found by BRBeans, then use either of the following to mark the Rule:

- Use the `setReady(false)` method in the Rule Management APIs
- Use the Rule Management Application to mark the rule as not ready

## Rule results

In general, a rule can return any type of result that makes sense for the business purpose of the rule. The return type on the `fire()` method is `java.lang.Object` so any Java object can be returned, including arrays. You cannot return a Java primitive since the results must be an object. However, you can return the object form of the primitives. For example, you can return a `java.lang.Integer` instead of an `int`. If the rule is fired remotely, the returned value must implement `java.io.Serializable`.

## Dependent rules

When a business rule triggers other business rules as part of its implementation, the rules that are triggered are called **dependent rules** of the first rule. An example is the `RuleAND` rule implementor supplied with Business Rule Beans (BRBeans). It uses two or more dependent rules, each of which is assumed to return a true or false value. When a rule with `RuleAND` as its implementor is triggered, it triggers each of its dependent rules and a logical AND operation is performed on all of the returned results. The result of this AND operation is returned as the result of the top-level rule.

Dependent rules are specified in the attributes of the top-level rule where the fully qualified name of each dependent rule is listed. When the top-level rule is triggered, an array of dependent rule names is passed to the rule implementor's `init()` method. They are stored here until they are triggered by the `fire()` method.

**Note:** The BRBeans framework does not ensure that the dependent rules specified in the enterprise beans are actually triggered. Triggering the dependent rules and interpreting their results is entirely up to the rule implementor of the top-level rule.

Dependent rules can be nested within other dependent rules. In other words, a dependent rule of some particular rule can have its own dependent rules which, in turn, can have their own dependent rule and so on. The BRBeans framework does not place any restriction on the number of levels that dependent rules can be nested. The only practical restriction is the complexity of the rule set that is built up when dependent rules are nested many levels deep.

## BRBeans run-time environment

The Business Rule Beans (BRBeans) run-time environment is used to find and trigger rules.

The BRBeans run-time environment is made up of two parts:

- Code that runs on the **client** ("client" here meaning wherever the trigger point is located). This consists of code that does the following:
  - Finds the specified rules
  - Decides where the rules should be triggered
  - Calls the fire method on all rules
  - Combines the results from the rules
- Code that runs on the **server**. This consists of enterprise beans used to represent rules and rule folders. These enterprise beans do the following:
  - Provides for business rule persistence
  - Provides query functions that the client part of the run time can use to find rules to be triggered

## BRBeans run-time behavior

Business Rule Beans (BRBeans) run-time behavior can best be described by giving a simple example of a trigger point selecting, executing, and then responding to the results of a business rule.

The first step in triggering a rule is for the trigger point framework to perform a query on the rule server to determine which rules to trigger. The main item used for the query is the fully qualified rule name. Other items used in the query include the start and end date, whether this is a classifier, the classification of the rule, and whether the rule is marked "ready". This query returns zero or more rules. If there is at least one rule, the trigger point assembles the data that is sent as parameters to each rule. The trigger point then loops through the list of rules invoking the fire() method on each and passing the parameters. The results are combined depending on the combining strategy used.

When the trigger point framework invokes fire on a rule, it instantiates the RuleImplementor and uses it to do the actual work (to execute the rule algorithm or test). Once it has arrived at a result, the RuleImplementor returns that result. For constraint rules (ones that arrive at a boolean true or false answer) the returned value is, by convention, a ConstraintReturn. A ConstraintReturn is a data structure that indicates whether the constraint was satisfied. If not, the ConstraintReturn indicates what went wrong. For derivation rules (ones that calculate a single, generally non-boolean value), the return value can be any type. In the simplest case, the return value from each RuleImplementor is returned back to the trigger point where it is analyzed to determine what action to take.

The following is an overview of what happens when the maxTruckDriverHours rule is triggered:

A rule exists named maxTruckDriverHours. The purpose of this rule is to check that the number of hours entered by a user for a particular truck driver does not exceed the maximum allowed value. This rule contains an initialization parameter list consisting of a single value of 8. This rule is bound to a RuleImplementor class called MaxRuleImpl. MaxRuleImpl tests the parameter it is passed against the initialization list value and returns a ConstraintReturn. The ConstraintReturn is set to true if the passed parameter is less than or equal to the initialization value.

Otherwise, a `ConstraintReturn` is set to false and some information is added that describes which values were compared and why the test failed.

When this rule is triggered, the following details the trigger point process:

1. During the execution of the application, the application reaches a point where it needs to verify that the number of truck driver hours that was entered is valid. The application code invokes a simple trigger point passing the name of the rule to be triggered and a parameter list containing the entered hours for the driver.
2. The trigger point framework performs a query on the rule server to find the rule with the specified name. It receives back a sequence of rule objects. In this case, this sequence contains one rule, `maxTruckDriverHours`.
3. The framework determines whether this rule is to be triggered on a local or remote machine. If local, the framework gets a local copy of the rule object and calls the `fire` method on the copy. If remote, the framework calls the `fire` method on the enterprise bean reference. The parameter list containing the weight is passed to the `fire` method.
4. The `maxTruckDriverHours` rule (either the copy or the enterprise bean itself) creates an instance of the rule implementor class, `maxRuleImpl`, if it does not already have one. When a new rule implementor instance is created, the rule calls its `init` method passing any initialization parameters defined for the rule. In this case, the initialization parameter list contains the single value 8. If the rule already has a rule implementor instance, it uses that one and does not call the `init` method.
5. The `maxTruckDriverHours` rule calls the `fire` method on the rule implementor instance. The firing parameters passed to the trigger point are passed to the rule implementor and are possibly modified by any firing parameters defined in the rule itself. In this case, the firing parameters are passed directly from the trigger point.
6. The `maxRuleImpl` returns a `ConstraintReturn` object to the rule that indicates the result of its comparison. This `ConstraintReturn` is returned to the trigger point framework and ultimately to the application.
7. The application checks the value in the `ConstraintReturn` and takes the appropriate action.

## BRBeans run-time exception handling

Business Rule Beans (BRBeans) defines one general exception class for exceptions that might be exposed to the user. All other BRBeans exceptions inherit from this class. The name of this class is `com.ibm.websphere.brb.BusinessRuleBeansException`. A `BusinessRuleBeansException` is generally thrown when an unexpected error occurs within BRBeans. A `BusinessRuleBeansException` might have information in it about the original exception that caused the error. Doing a `printStackTrace` on the `BusinessRuleBeansException` prints out this information and the stack trace for the `BusinessRuleBeansException` itself. Also, there are methods to access the original exception programmatically, if desired.

BRBeans also defines a **`ConstraintViolationException`**, which extends `BusinessRuleBeansException`. A `ConstraintViolationException` is thrown if the `ThrowViolationCombiningStrategy` is specified on the `TriggerPoint` and the rule returns a false value (either a `ConstraintReturn` or a boolean).

Finally, BRBeans defines two exceptions, **`NoRuleFoundException`** and **`MultipleRulesFoundException`**, that are thrown by some of the predefined filtering

strategies if an unexpected number of rules is found on a trigger point call. These two exceptions both extend `UnexpectedRulesFoundException` which, in turn, extends `BusinessRuleBeansException`.

## Rule implementors

A rule implementor, in terms of Business Rule Beans, is an algorithm written in Java that implements the `BRBeans RuleImplementor` interface.

A Business Rule Beans (BRBeans) Rule is a persistent object that exists on the BRBeans Rule server. One of the rule's persistent attributes, in addition to `startDate`, `endDate`, `initParams`, and so on, is `javaRuleImplementorName`, which is the name of its rule implementor.

BRBeans supplies a number of predefined rule implementor classes that can be used in user-defined BRBeans rules (see the BRBeans Javadoc ([../javadoc/ee/com/ibm/websphere/brb/RuleImplementor.html](http://../javadoc/ee/com/ibm/websphere/brb/RuleImplementor.html))) to implement the `BRBeans RuleImplementor` Interface. The Java source code for these rule implementors is supplied as BRBeans sample code in the `com.ibm.websphere.brb.implementor` package. This sample code, packaged in a JAR file, appears in the CLASSPATH of the BRBeans Rule Server (for "remote" firing) or is co-located in the CLASSPATH of the application or applications using it (for "local" firing). Typically, the `RuleImplementor` is in the application EAR file.

Using standard Java development tools, you can externalize BRBeans by attaching them to either enterprise beans or ordinary Java objects. Programming a new rule implementation in Java is typically a simple process. If you write your own rule implementor, you must create a new Java class that implements the `com.ibm.websphere.brb.RuleImplementor` interface. This class must implement the following methods:

### Default constructor

The class must have a default, no-argument constructor so that it can be instantiated when a rule using it is triggered.

**init** The `init` method comes from the `RuleImplementor` interface and is called when the rule implementor is first. Its purpose is to perform an initialization needed by the rule implementor instance before it is actually fired. The following parameters are passed to the `init` method:

#### The initialization parameters defined for the rule being triggered

These can be any parameters needed to properly initialize the rule implementor instance. Often the initialization parameters consist of constants required by the algorithm. For example, when using a rule implementor that checks whether a number is greater than a threshold value, the threshold value normally is passed as an initialization parameter. This parameter is null if there are no initialization parameters for the rule.

#### An array of names of dependent rules for the rule being triggered

Normally, the rule implementor stores these names to be used when the `fire` method is called. These dependent rules are intended to be triggered as part of the algorithm performed by the rule implementor. See `Dependent Rules` for more information. This parameter is null if there are no dependent rules defined for the rule.

#### The user-defined data for the rule being triggered

This data is completely defined by the user of the Business Rule

Beans (BRBeans). BRBeans does not interpret this data in any way. This parameter is null if there is no user-defined data defined for the rule.

**A reference to the actual rule being triggered**

This can be used to extract attribute values from the rule, if needed.

**fire** The fire method comes from the RuleImplementor interface. This method is called to perform the algorithm of the rule implementor. Any desired algorithm can be performed here. Normally, a value is returned by the fire method that is ultimately returned as the result of triggering the rule. The following parameters are passed to the fire method:

**The TriggerPoint object that is being used to trigger the rule**

This parameter is needed if the rule has dependent rules that the fire method needs to trigger.

**The target object for this particular trigger call**

This parameter can be any object that is thought of as the target of the rule. However, the parameter can be null.

**A reference to the actual rule being triggered**

This parameter can be used to extract attribute values from the rule, if needed.

**The firing parameters for this particular trigger call**

Normally, these parameters are the firing parameters passed by the code that invoke the trigger point. However, these can be overridden by specifying firing parameters on the rule itself. Wherever they ultimately come from, these are the parameters that the rule implementor needs at run time to perform its function. Normally, these are run-time variables that are to be processed by the rule implementor. For example, when using a rule implementor that checks whether a number is greater than a threshold value, the number to be checked normally is passed as a firing parameter. This parameter is null if no firing parameters are passed by the caller and none are defined on the rule itself.

**getDescription**

getDescription comes from the RuleImplementor interface. The purpose of this method is to return a text string that describes the function of the rule implementor. This information might be displayed on a user interface to help select what implementor to use. This method, however, is currently not used by the BRBeans framework. Users can incorporate this information if they create their own rule implementor. For additional information, see the BRBeans Javadoc ([../javadoc/ee/com/ibm/websphere/brb/RuleImplementor.html](http://../javadoc/ee/com/ibm/websphere/brb/RuleImplementor.html)) for the RuleImplementor interface.

## Trigger point framework

A trigger point is the location in a method of an object where externalized business rules are invoked.

The proper placement of trigger points can add substantially to the flexibility and speed with which a business application adapts to new business practices.



Wherever a trigger point is placed in user-written code, the Business Rule Beans (BRBeans) trigger point framework needs to do the following:

1. Assemble the parameter list to send to the rules
2. Find the potential rules that apply
3. Filter out any rules which do not apply (optional)
4. Fire the rules in the filtered rule set
5. Combine the results of the rule firings in some meaningful way

The application code that contains the trigger point needs to perform the following functions:

1. Establish a value for the target object. Usually the target object is the object in which the trigger point is encountered. The target object is one of the parameters passed to the fire method of the RuleImplementor. However, this is an optional parameter. If the rule implementor does require a target object, null can be passed instead.
2. Build the array of objects containing the run-time parameters needed to satisfy the trigger point's business purpose. This array is normally passed as one of the parameters of the fire method of the RuleImplementor. If firing parameters are specified on the rule itself, then those firing parameters are passed instead of the ones passed by the caller.
3. Invoke the trigger(), triggerClassifier(), or triggerSituational() method of the TriggerPoint class.
4. Catch and handle any exceptions that might occur as a result of firing the rules. Otherwise, take action based upon the rule firing results.

The two simple trigger methods, trigger and triggerClassifier, perform their function in four steps:

1. Find the rules
2. Filter out those rules which are not desired
3. Fire the remaining rules
4. Combine the results and return to the caller

The complex trigger method, triggerSituational does this sequence of steps twice. In the first phase, the method performs the four steps once to find a rule that returns a classification. This classification is fed into the second phase. The second phase triggers rules that have the name specified in the triggerSituational method and have a classification equal to the value returned by the first phase.

How each of these steps is performed can be modified through various methods on the TriggerPoint object. The implementation of each step is defined by a **strategy** object. For more information on strategies, see Administering strategy objects to control triggers.

## Trigger points

Examples of how to code a trigger point call are provided in the following topics:

- Simple trigger point
- Classifier trigger point
- Situational trigger point

### Simple trigger point

A simple trigger point is used to trigger a rule or rules specified by name. This type of trigger point is used by invoking the trigger method on an instance of the

TriggerPoint class. All rules with the specified name are triggered and the results are combined using the CombiningStrategy specified on the TriggerPoint object. This type of trigger point only finds rules that are not marked as classifiers.

The following shows an example of using a simple trigger point to trigger a rule named isSeniorCitizen (in the com/acme/ageRules folder), which determines whether a person is classified as a senior citizen based on the passed in age:

```
...
// create an instance of TriggerPoint for triggering the rule and specify that the
// ReturnFirstCombiningStrategy is to be used to return only the first result if
// multiple rules are found.
TriggerPoint tp = new TriggerPoint();
tp.setCombiningStrategy(CombiningStrategy.RETURN_FIRST, TriggerPoint.ALL_RULES);

// define parameter list that's passed to the rule
Object [ ] plist = new Object[1];

// define age of person to be tested
Integer age = new Integer(64);

// define name of rule to be fired
String ruleName = "com/acme/ageRules/isSeniorCitizen";

// define result of rule firing
Object result = null;

// initialize parameter list
plist[0] = age;

try {

// fire "com/acme/ageRules/isSeniorCitizen" rule passing parameter list containing
// age.
// Note: in this case the target object is not used and could be null.
    result = tp.trigger(this, plist, ruleName);

// put result into usable format. A single result is returned since we specified to
// use the ReturnFirstCombiningStrategy. By default an array of results is returned.
    boolean seniorCitizen = ((Boolean)result).booleanValue();

// make use of result
    if( seniorCitizen ) {
        ...
    }
}
catch(BusinessRuleBeansException e ) {

// handle exception
    ...
}
}
```

### **Classifier trigger point**

A classifier trigger point is identical to a simple trigger point except that it only finds rules marked as classifiers. Classifiers are rules that determine what sort of business situation is present. These rules then return a classification string that indicates the result.

Usually these rules are used as part of a situational trigger point, but they also can be triggered on their own. This type of trigger point is used by invoking the triggerClassifier method on an instance of the TriggerPoint class.

The following shows an example of using a classifier trigger point to trigger a rule named `determineCustomerLevel` (in folder `com/acme/customerClassifiers`). This rule classifies customers into levels (gold, silver, and bronze) based on their spending history.

```
...
// create an instance of TriggerPoint for triggering the rule and specify that the
// ReturnFirstCombiningStrategy is to be used to return only the first result if
// multiple rules are found.
TriggerPoint tp = new TriggerPoint();
tp.setCombiningStrategy(CombiningStrategy.RETURN_FIRST, TriggerPoint.ALL_RULES);

// define parameter list that's passed to the rule
Object [ ] plist = new Object[1];

// information about the customer to be checked is stored in this object
Customer cust = ...;

// define name of rule to be fired
String ruleName = "com/acme/customerClassifiers/determineCustomerLevel";

// define result of rule firing
Object result = null;

// initialize parameter list
plist[0] = cust;

try {

    // fire "com/acme/customerClassifiers/determineCustomerLevel" rule passing
    // parameter list containing the customer to be checked.
    // Note: in this case the target object is not used and could be null.
    result = tp.triggerClassifier(this, plist, ruleName);

    // put result into usable format. A single result is returned since we
    // specified to use the ReturnFirstCombiningStrategy. By default an array of
    // results would be returned.
    String customerLevel = (String) result;

    // make use of result
    if( customerLevel.equals("Gold") ) {
        ...
    } else if ( customerLevel.equals("Silver") ) {
        ...
    } else if ( customerLevel.equals("Bronze") ) {
        ...
    } else {
        ...
    }
}
catch(BusinessRuleBeansException e ) {

    // handle exception
    ...

}
```

### Situational trigger point

A situational trigger point is used when the rule or rules to be triggered depend on the business situation.

This example evaluates a customer's past purchasing history to place them into one of three levels: Gold, Silver, or Bronze. Their classification determines how much of a discount they receive.

To use a situational trigger point to handle this case, it is first necessary to define four rules:

- one **classifier** rule to determine under which of the three levels to classify the customer
- three **classified** rules to determine the actual discount to offer

All of the classified rules have the same name and are marked as applying to one of the three customer levels by specifying the level in its classification attribute. For example, the rule to determine the discount for a Gold level customer will contain the string "Gold" in its classification attribute.

The situational trigger point takes two rule names as input: the name of the classifier rule and the name of the classified rule. The situational trigger point then proceeds in two phases:

1. Find the specified classifier rule and trigger it to generate a classification string.
2. Find the rules that have the name specified for the classified rule and have a classification attribute equal to the classification string returned by the first phase.

These rules then are triggered to produce the final result, in this case the discount to offer.

The following shows an example of a situational trigger point used to handle the case described previously:

```
...
// create an instance of TriggerPoint for triggering the rule and specify that the
// ReturnFirstCombiningStrategy is to be used to return only the first result if
// multiple rules are found.
TriggerPoint tp = new TriggerPoint();
tp.setCombiningStrategy(CombiningStrategy.RETURN_FIRST, TriggerPoint.ALL_RULES);

// define parameter list that's passed to the classifier rule
Object [ ] classifierPlist = new Object[1];

// define parameter list that's passed to the classified rule
Object [ ] classifiedPlist = new Object[1];

// information about the customer to be checked is stored in this object
Customer cust = ...;

// define name of classifier rule to be fired
String classifierRuleName = "com/acme/customerClassifiers/determineCustomerLevel";

// define name of classified rule to be fired
String classifiedRuleName = "com/acme/discountRules/determineDiscount";

// define result of rule firing
Object result = null;

// initialize parameter lists
classifierPlist[0] = cust;
classifiedPlist[0] = cust;

try {
    // fire the rules to get the discount to offer
    // Note: in this case the target object is not used and could be null.
    result = tp.triggerSituational(this, classifiedPlist, classifierPlist,
        classifiedRuleName, classifierRuleName);

    // put result into usable format. A single result is returned since we
    // specified to use the ReturnFirstCombiningStrategy. By default an array
```

```

    // of results would be returned.
    Float discountToOffer = (Float) result;

    // make use of result
    ...
}
catch(BusinessRuleBeansException e ) {

    // handle exception
    ...
}

```

## As Of Date

An "As Of Date" can cause rules to be triggered as if the given date is the current date. This is especially useful when you want to test a rule, see what effect a future change in rules or regulations may have on the overall framework, or see what past or future rates, discounts, or both might be.

Normally, a rule only can be triggered if it is "in effect" (see "Rule States") as of the current date and time.

To set an "As Of Date", call the `setAsOfDate()` method on the `TriggerPoint` object and pass the date that you want to be used. To use the current date again, call `unsetAsOfDate` or `setAsOfDate` and pass null for the date.

## Predefined strategy objects

The following is a list of predefined strategy objects that are provided in Business Rule Beans:

### **FindingStrategy**

Accesses the data store and returns those rules that meet the search criteria specified

### **FilteringStrategy**

Takes the list of rules that were found by the `FindingStrategy` and filters out the rules that should not be fired

### **FiringStrategy**

Takes the rules that were found by the `FindingStrategy`, (possibly modified by the `FilteringStrategy`), fires them each in order, and returns an array containing the results of each rule

### **CombiningStrategy**

Takes the results of the rules that are fired by the `FiringStrategy` and combines them to form a reasonable result to the `TriggerPoint` caller.

### **FindingStrategy method**

The job of the `FindingStrategy` is to access the data store and return those rules that meet the search criteria specified. There are two `FindingStrategy` classes provided by Business Rule Beans (BRBeans):

- `DefaultClassifierFindingStrategy`
- `DefaultNonClassifierFindingStrategy`

Both of these strategies perform a case-sensitive search for Rules that are marked "ready" and match the given search criteria. Results are ordered by precedence from highest to lowest (the first rule in the array has the numerically smallest precedence, the next rule has the next smallest precedence, and so on). If no rules

are found, then an empty array is returned. The former strategy returns classifier rules (`classifier=true`) only and the latter returns non-classifier rules (`classifier=false`) only.

These default strategies are used automatically by the `TriggerPoint`. There is no need to call `setFindingStrategy` to use these strategies. Instances of these two default finding strategies are stored in static constants defined on the `FindingStrategy` interface.

### FilteringStrategy method

The job of the `FilteringStrategy` is to take the list of rules that were found by the `FindingStrategy` and filter out the rules that should not be fired. There are three sets of filtering strategies used in `TriggerPoint`:

- strategy for **zero** rules found
- strategy for **one** rule found
- strategy for **multiple** rules found

A different strategy can be used for each of these scenarios, along with different strategies for classifier and non-classifier rules. The zero rules strategy is invoked if no rules are found by the finding strategy, the one rule strategy is invoked if exactly one rule is found and the multiple rules strategy is invoked if more than one rule is found.

Business Rule Beans (BRBeans) provides the following filtering strategies that can be used:

#### Accept Any

BRBeans utilizes all of the rules found (this is the default).

#### Accept One

BRBeans expects one rule only.

#### Accept First

BRBeans utilizes the first rule found.

#### Accept Last

BRBeans utilizes the last rule found.

Instances of these filtering strategies are stored in static constants defined in the `FilteringStrategy` interface. You can use these for setting the strategies on a `TriggerPoint`.

As an example, here is one common way to use filtering strategies. You want to ensure that exactly one rule is found on a `TriggerPoint` call. Thus, set all three strategies (zero rules, one rule, and multiple rules) for this `TriggerPoint` to `FilteringStrategy.ACCEPT_ONE`. This strategy throws an exception if the number of rules is not exactly one. The following sequence of method calls accomplishes this for `TriggerPoint tp`:

```
tp.setNoRulesFilteringStrategy(FilteringStrategy.ACCEPT_ONE, TriggerPoint.ALL_RULES);
tp.setOneRuleFilteringStrategy(FilteringStrategy.ACCEPT_ONE, TriggerPoint.ALL_RULES);
tp.setMultipleRulesFilteringStrategy(FilteringStrategy.ACCEPT_ONE,
                                     TriggerPoint.ALL_RULES);
```

### FiringStrategy method

The `FiringStrategy` takes the rules that were found by the `FindingStrategy`, (possibly modified by the `FilteringStrategy`), fires them each in order, and returns an array containing the results of each rule.

A single default `FiringStrategy` is provided by Business Rule Beans (BRBeans) as all of the rules are fired in the same way. This implementation takes each rule in order and performs the following steps:

1. Determines what firing parameters to pass to the rule. If there are no firing parameters specified for this rule, the implementation uses the firing parameters passed on the `TriggerPoint` call. Otherwise, it uses the firing parameters specified in the rule in place of the parameters passed on the `TriggerPoint` call.
2. Calls the `fire` method on the rule and passes the firing parameters from the first step.

Unexpected exceptions result in a `BusinessRuleBeansException` being thrown that contains the original exception.

### **CombiningStrategy method**

The job of the `CombiningStrategy` is to take the results of the rules that are fired by the `FiringStrategy` and combine them to form a reasonable result to the `TriggerPoint` caller. Business Rule Beans (BRBeans) provides several combining strategies to be used in applications:

#### **Return All**

Returns the results from all of the rules fired in an array (this is the default)

#### **Return First**

Returns only the result from the first rule fired

#### **Return Last**

Returns only the result from the last rule fired

#### **Return AND**

Returns the logical AND of the results from all the rules fired. This strategy requires that all of the results returned by the fired rules are either `ConstraintReturn` objects or `java.lang.Boolean` objects. An exception is thrown if this is not the case.

#### **Return OR**

Returns the logical OR of the results from all of the rules fired. This strategy requires that all of the results returned by the fired rules are either `ConstraintReturn` objects or `java.lang.Boolean` objects. An exception is thrown if this is not the case.

#### **Throw Violation**

Throws a `ConstraintViolationException` containing all of the failed `ConstraintReturn` objects if any `ConstraintReturn` objects contain `false`. Otherwise, it returns a `true ConstraintReturn`.

Instances of these combining strategies are stored in static constants defined in the `CombiningStrategy` interface. You can use these for setting the strategies on a `TriggerPoint`. For example, the following method call sets the `CombiningStrategy` on `TriggerPoint tp` to be the `Return_First` strategy:

```
tp.setCombiningStrategy(CombiningStrategy.RETURN_FIRST, TriggerPoint.ALL_RULES);
```

## **Customized strategy objects**

The process of triggering a rule or set of rules is controlled by a set of strategy objects. The following four strategies are used each time a rule is triggered:

### **FindingStrategy**

The `FindingStrategy` accesses the persistent data store to find the set of

rules matching the search criteria passed to the trigger call. The search criteria are based on the rule ID information passed on the trigger call. The set of rules found is passed to the `FilteringStrategy`.

### **FilteringStrategy**

The `FilteringStrategy` can change the set of rules that were found by the `FindingStrategy`. The set of rules returned is the set that are fired by the `FiringStrategy`.

### **FiringStrategy**

The `FiringStrategy` fires the rules found by the `FindingStrategy`, which may be modified by the `FilteringStrategy`. It gathers the results of the individual rules and passes them to the `CombiningStrategy`.

### **CombiningStrategy**

The `CombiningStrategy` takes the results from firing the rules and combines them to produce the final result of the trigger.

Each `TriggerPoint` object has its own set of strategies that can be changed independent of any other `TriggerPoint` object. There is a set of default strategies that are used by the `TriggerPoint` if none are explicitly set.

For each of the four strategies, you can set different strategies for **classifier** rules and for **non-classifier** rules. The strategies set for classifier rules are used when the Business Rule Beans (BRBeans) framework is triggering a classifier rule. The strategies for non-classifier rules are used in all other cases.

It is also possible to set three different sets of filtering strategies:

- one to be used if no rules are found
- one to be used if exactly one rule is found
- one to be used if more than one rule is found

This capability can be used to set up filtering strategies that throw exceptions if the expected number of rules is not found.

Strategy classes must implement one of the strategy interfaces provided by BRBeans in the `com.ibm.websphere.brb` package:

- `FindingStrategy`
- `FilteringStrategy`
- `FiringStrategy`
- `CombiningStrategy`

Users can write their own strategy implementations to perform special functions not performed by the predefined implementations. Write these strategy implementations with care since part of the functionality of the BRBeans framework is replaced when you write a custom strategy. One simple example of writing a custom strategy is creating a new firing strategy that logs every rule that is fired.

The basic requirement for a strategy implementation is that it implements the appropriate strategy interface.

For the **filtering** and **combining** strategies, create a class that implements either `FilteringStrategy` or `CombiningStrategy` and either the `filterRules()` method (for `FilteringStrategy`) or the `combineResults()` method (for `CombiningStrategy`) to perform the required functions. At run time, create an instance of the new class



and pass it to the `TriggerPoint` object using the appropriate set method so that the new strategy is used when rules are triggered using that `TriggerPoint`.

The **finding** and **firing** strategies are more complicated to customize since they provide more function than the simple filtering and combining strategies. Default finding and firing strategy implementations are provided that define a general outline of the steps necessary to perform the function. It is suggested that you subclass these when you customize your own strategies and then override the desired methods on the default implementation to provide the new behavior.

The BRBeans Javadoc provides more information about the `FindingStrategy` ([../javadoc/ee/com/ibm/websphere/brb/FindingStrategy.html](http://../javadoc/ee/com/ibm/websphere/brb/FindingStrategy.html)), `FilteringStrategy` ([../javadoc/ee/com/ibm/websphere/brb/FilteringStrategy.html](http://../javadoc/ee/com/ibm/websphere/brb/FilteringStrategy.html)), `FiringStrategy` ([../javadoc/ee/com/ibm/websphere/brb/FiringStrategy.html](http://../javadoc/ee/com/ibm/websphere/brb/FiringStrategy.html)), and the `CombiningStrategy` ([../javadoc/ee/com/ibm/websphere/brb/CombiningStrategy.html](http://../javadoc/ee/com/ibm/websphere/brb/CombiningStrategy.html)) interfaces.

## Customized rule implementors

To write your own rule implementor, create a new Java class that implements the `com.ibm.websphere.br.RuleImplementor` interface. This class must implement the following methods:

### Default constructor

The class must have a default, no-argument constructor so that it can be instantiated when a rule using it is triggered.

**init** The `init` method comes from the `RuleImplementor` interface and is called when the rule implementor is first created. Its purpose is to perform an initialization needed by the rule implementor instance before it is actually fired. The following parameters are passed to the `init` method:

### The initialization parameters defined for the rule being triggered

These can be any parameters needed to properly initialize the rule implementor instance. Often the initialization parameters consist of constants required by the algorithm. For example, when using a rule implementor that checks whether a number is greater than a threshold value, the threshold value normally is passed as an initialization parameter. This parameter is null if there are no initialization parameters for the rule.

### An array of names of dependent rules for the rule being triggered

Normally, the rule implementor stores these names to be used when the `fire` method is called. These dependent rules are intended to be triggered as part of the algorithm performed by the rule implementor. See *Dependent Rules* for more information. This parameter is null if there are no dependent rules defined for the rule.

### The user-defined data for the rule being triggered

This data is completely defined by the user of the Business Rule Beans (BRBeans). BRBeans does not interpret this data in any way. This parameter is null if there is no user-defined data defined for the rule.

### A reference to the actual rule being triggered

This can be used to extract attribute values from the rule, if needed.

**fire** The fire method comes from the RuleImplementor interface. This method is called to perform the algorithm of the rule implementor. Any desired algorithm can be performed here. Normally, a value is returned by the fire method that is ultimately returned as the result of triggering the rule. The following parameters are passed to the fire method:

**The TriggerPoint object that is being used to trigger the rule**

This parameter is needed if the rule has dependent rules that the fire method needs to trigger.

**The target object for this particular trigger call**

This parameter can be any object that is thought of as the target of the rule. However, the parameter can be null.

**A reference to the actual rule being triggered**

This parameter can be used to extract attribute values from the rule, if needed.

**The firing parameters for this particular trigger call**

Normally, these parameters are the firing parameters passed by the code invoking the trigger point. However, these can be overridden by specifying firing parameters on the rule itself. Wherever they ultimately come from, these are the parameters that the rule implementor needs at run time to perform its function. Normally, these will be run-time variables that are to be processed by the rule implementor. For example, when using a rule implementor that checks whether a number is greater than a threshold value, the number to be checked normally is passed as a firing parameter. This parameter is null if no firing parameters are passed by the caller and none are defined on the rule itself.

**getDescription**

getDescription comes from the RuleImplementor interface. The purpose of this method is to return a text string that describes the function of the rule implementor. This information might be displayed on a user interface to help a user select what implementor to use. This method, however, is currently not used by the BRBeans framework. For additional information, see the RuleImplementor ([../javadoc/ee/com/ibm/websphere/brb/RuleImplementor.html](http://../javadoc/ee/com/ibm/websphere/brb/RuleImplementor.html)) interface in the BRBeans Javadoc.

## Rule management command

The Rule management command assists the user in performing high-level administration of rules and rule folders.

This includes the capability to create, modify, delete, import, or export rules or rule folders. This command can be used initially by the programmer to define rules interactively and then used by the domain analyst for rule management tasks. You can use the following files:

- On Microsoft Windows platforms, rulemgmt.bat
- On Unix platforms, rulemgmt.sh

**Syntax**

rulemgmt properties-file [host-address port-number]

### Parameters

#### <properties-file>

The fully qualified name of a file containing the JNDI names of the rule EJBs for the rule set that is to be accessed.

The following must be specified in the file:

```
RuleJndi=<JNDI name of the Rule EJB>
RuleFolderJndi=<JNDI name of the RuleFolder EJB>
RuleHelperJndi=<JNDI name of the RuleHelper EJB>
```

See <WAS\_HOME>/bin/brbeansDefaultProperties for an example.

### Arguments

#### host-address

This is the host name of the name server. The default is the local host.

#### port-number

This is the port number of the name server. The default is 2809.

## Rule importer command

The rule importer command imports rules into a database from one or more XML documents.

The rule importer command can be invoked using the Rule Management Application (RMA). The user interface in RMA provides some assistance in specifying the parameters required by the importer. Alternatively, the rule importer can be invoked from the command line using the following scripts:

- For Microsoft Windows platforms, ruleimporter.bat
- For UNIX platforms, ruleimporter.sh

### Syntax

```
ruleimporter <properties-file> <import-files> [options]
```

### Parameters

#### <properties-file>

The fully qualified name of a file containing the JNDI names of the Business Rule Beans (BRBeans) enterprise beans for the rule set that is to be accessed. Refer to BRBeans Properties file for a definition of the contents of this file. This parameter is required.

#### <import-files>

One or more fully qualified names of the files containing XML rule definitions to be imported. These files must contain XML in the format defined in <WAS\_HOME>\bin\brb.dtd file. This parameter is required.

### Options

#### -[v]erbose

Shows verbose output while importing. This shows the rule definition of every rule that is imported.

#### -[t]est

Parses the input files only and does not create rules on the application server. This will ensure that there are no errors in the syntax of the rule

definitions provided in the XML document. Combined with the **-verbose** option, it also can be used to see exactly what rules will be imported.

**-[u]pdate**

Update the existing rule with values from the input file when a rule in an input file has the same primary key as an existing rule. If this option is not specified, then any rule with the same primary key as an existing rule causes an error and that rule is not imported.

**-[c]ommiteach**

Performs a commit after each rule is created rather than creating all of the rules in a single transaction. If this option is not specified, then all rules are created in a single transaction. This means that if any rule causes an error, the entire transaction is rolled back and none of the rules are imported. If **-commiteach** is specified and a rule causes an error, only that rule is not imported. Other rules are still imported.

**-[h]ost <host-name>**

Specifies the name of the host for the name server. The default is the local host.

**-[p]ort <port-number>**

Specifies the port number for the name server. The default is 2809.

## Rule exporter command

The rule exporter command exports rules from a database into an XML document.

The rules that are exported are determined by an XML document, which is provided to the command. The rule exporter function can be invoked using the Rule Management Application (RMA). The user interface in RMA provides some assistance in specifying the parameters required by the exporter. Alternatively, the rule exporter can be invoked from the command line using the following scripts:

- For Microsoft Windows platforms, ruleexporter.bat
- For UNIX platforms, ruleexporter.sh

### Syntax

```
ruleexporter <properties-file> <export-list-files> [options]
```

### Parameters

**<properties-file>**

The fully qualified name of a file containing the JNDI names of the BRBeans enterprise beans for the rule set that is to be accessed. Refer to BRBeans Properties file for a definition of the contents of this file. This parameter is required.

**<export-list-files>**

One or more fully qualified names of files containing a list of rules, folders, or both to be exported. These files must contain XML in the format defined in the <WAS\_HOME>\AppServer\bin\brb-export-list.dtd file. This parameter is required.

### Options

**-[v]erbose**

Shows verbose output while exporting.

**-[o]utput <file-name>**

Specifies the name of the output file where the XML rule definitions are stored. This is a required parameter.

**-[h]ost <host-name>**

Specifies the name of the host for the name server. The default is the local host.

**-[p]ort <port-number>**

Specifies the port number for the name server. The default is 2809.

## BRBeans properties file

Applications that use the Business Rule Beans (BRBeans) enterprise beans (this includes applications that trigger rules or use the rule management APIs) must specify the JNDI names for these enterprise beans so that the application can find them at run time. If the application is running in a J2EE client container, in a servlet, or on the application server itself (for example, as part of another enterprise bean), then these names probably have been specified by the person who configured the application. If the application is not running in a container, the names must be specified some other way. The BRBeans properties file provides a way to do this.

At run time, the BRBeans code looks for a special Java property that identifies the name of the properties file. This Java property can be specified on the command line as `-DbrbPropertiesFile=<file_name>`. The file specified is expected to contain the JNDI names used to find the BRBeans enterprise beans. The BRBeans framework uses these names when it needs to locate the enterprise beans.

When an application attempts to reference BRBeans enterprise beans, the code first looks for the `brbPropertiesFile` Java property. If this property is specified, the names listed in that file are used to find the enterprise beans and to override any EJB references that were specified in the container (if the application is running in a container). If the property is not specified, then BRBeans attempts to use the EJB references specified in the container.

The host name and port number used to access the name server also can be set in this file. If these are not specified, the BRBeans framework uses the name server used by the container in which the application is running. If the application is not running in a container, then `localhost` is used for the host name and `2809` is used for the port number.

The properties file must be in the following format (entries can be specified in any order):

```
host=<host-name-for-name-server>
port=<port-number-for-name-server>
RuleJndi=<JNDI-name-for-Rule-EJB>
RuleFolderJndi=<JNDI-name-for-RuleFolder-EJB>
RuleHelperJndi=<JNDI-name-for-RuleHelper-EJB>
```

### Location

A default properties file is shipped as  
`<WAS_HOME>\AppServer\bin\brbeansDefaultProperties.`

There are a set of JAR files that conform to the BRBeans<database-type>.jar naming convention (depending on the type of database that you want to use). If the JAR files are used without changing the JNDI names, then you also can use the default properties file.

#### **Usage note**

The file name still must be specified even if you want to use the default file. There is no file that is used automatically if the `brbPropertiesFile` property is not set.

The tools shipped with BRBeans (the Rule Management Application, the rule importer, and the rule exporter) all run outside of any container. Hence, the JNDI names need to be specified when these tools are run. The scripts for these tools all require that you pass a properties file name as a command line parameter. This name then is specified as the value for the `brbPropertiesFile` property when the tool is run.

## **Database considerations for BRBeans**

The following relational databases are supported by Business Rule Beans (BRBeans):

- IBM DB2
- IBM Cloudscape
- Microsoft SQL Server
- Oracle
- Sybase
- Informix

This documentation does not provide you with specific instructions on how to use any of these databases. For help with specific commands, consult the documentation that accompanied your database software. The following are general considerations for relational databases that are supported by BRBeans:

#### **Large character data**

There are several attributes in the BRBeans Rule enterprise beans that might contain large amounts of data. This includes fields such as: `businessIntent`, `dependentRules`, `description`, `firingParameters`, `initParameters`, `originalReq`, and `userDefinedData`. The value for these attributes is stored in a character type column within a database table. Whenever possible, the values are stored in large character fields like `LONG VARCHAR` (for DB2) and `TEXT` (for Sybase).

There are several cases where the use of large character fields is problematic, mostly in terms of a lack of query support. Refer to each of the supported database sections for details on the column type used for storing the values in these attributes.

#### **Isolation level**

All of the enterprise beans accessed in a transaction must specify the same isolation level. If your application contains enterprise beans that are used in the same transaction as the rules, you must do one of the following:

- Change the BRBeans enterprise beans (`Rule`, `RuleFolder`, and `RuleHelper`) to the same isolation level as your beans.

- Change your beans to the same isolation level as the BRBeans enterprise beans.
- Place the BRBeans enterprise beans in a different database than your enterprise beans and configure the application to run using the two-phase commit protocol. This causes the beans to run in different transactions; thereby removing the restriction that they need to have the same isolation level.

## Oracle considerations

### Large character data

The preferred Oracle data type for storing large character objects is CLOB. However, Oracle does not allow a CLOB to be queried. Because of this, a data type of VARCHAR2 is used by Business Rule Beans (BRBeans). A specific length must be specified when specifying VARCHAR2. The maximum length for a VARCHAR2 is 4000 bytes.

To determine the default size of VARCHAR2, look in the table.ddl file that was generated when you deployed the code. If the default size is not acceptable for your application, you can do one of the following:

#### Increase the size of the columns

Keep in mind that maximum size for a VARCHAR2 in Oracle is 4000. Increase the column size either by changing the value in the **create table statement** or by changing the schema mapping and deploying the BRBeans JAR file.

#### Change the schema mapping to specify CLOB

Do this for any of the attributes that you do not wish to query and then deploy the BRBeans JAR file.

### Isolation level

The default isolation level is REPEATABLE\_READ. Oracle does not support this isolation level. Therefore, the IBM WebSphere runtime environment converts this to the next highest isolation level, which in this case is SERIALIZED. Be aware that this isolation level tends to be overly restrictive as it prevents two clients from reading data at the same time. The BRBeansOracle.jar file specifies an isolation level of READ\_COMMITTED.

## Sybase considerations

### Allowing null values

By default, Sybase does not allow null values in string columns (like VARCHAR, TEXT, and so on). You can change this default value for a database using "isql" by issuing the following command:

```
sp_dboption databasename, "allow nulls by default", true
```

In this example "databasename" is your database name.

### Large character data

The large character data fields are stored in a column of type TEXT. Sybase allows TEXT fields to be queried only using the SQL "LIKE" operator. Queries against the columns that perform the SQL "IS NULL" or "IS NOT NULL" operations are not allowed by Sybase. The alternative is to specify a column type of VARCHAR. However, the maximum allowed size for a VARCHAR in Sybase is 255 characters. This is not considered a large enough value for storing firingParameters, initParameters, descriptions, and so on.

If performing "IS NULL" and "IS NOT NULL" type queries is important and the 255 character limitation is acceptable, change these column types to VARCHAR. To accomplish this, alter the schema mapping for the Rule bean and then deploy the BRBeans JAR file.

The query APIs (in the com.ibm.websphere.brb.query package) allow for "IS NULL" and "IS NOT NULL" type queries to be performed on several of these fields. In addition, the Rule Management Application allows the firing parameters to be queried in this manner. These queries fail on Sybase with the default column type of TEXT.

## **Informix considerations**

### **Large character data**

The preferred Informix data type for storing large character data is CLOB. However, Informix does not allow a CLOB to be queried. Because of this, a data type of LVARCHAR is used by BRBeans. The maximum length for an Informix LVARCHAR is 2,000 characters. If 2,000 characters is not acceptable and your application does not need to query these data types, you can change the schema mapping to specify CLOB. Then, deploy the BRBeans JAR file.

### **Custom properties for the data source**

When configuring the data source for your application, you must specify the following properties:

- ifxIFXHOST=Name of the physical machine on which the Informix instance is installed
- serverName=Informix instance name
- portNumber=Port number for which the Informix instance is configured
- informixLockModeWait=500

A setting of 500 causes a connection to wait for up to 500 seconds for a lock. If you have a busy system, this wait can appear to be a system hang. This setting has the same effect as running SET LOCK MODE TO WAIT 500 on the connection.

**Note:** The previous configuration values are subject to change. Consult your Informix documentation for updates.

## **Rule Management Application**

The Rule Management Application (RMA) is a tool that assists the user in performing high-level administration of rules and rule folders.

This includes the capability to create, modify, delete, import, or export rules or rule folders. The RMA tool can be used initially by the programmer to define rules interactively, and then by the domain analyst for rule management tasks. The main window for the RMA is the Rules Browser.

The column on the left side of the Rule Browser window shows a nested hierarchy of all of the existing rule folders. Click one of these folders to display the rules it contains. The names of these rules appear in the right column.

Navigate as you would in a typical file-management browser.

1. Click the "+" icon to expand by one level and click the "-" icon to collapse it.



2. Click a rule or folder name to highlight it, right-click the rule or folder name to launch a list of actions, or select an option from the main menu.

RMA is designed to be a general purpose tool for interactive management of rules. Alternatively, you can write your own user interface that is tailored more specifically to the domain in which you work. For instance, a domain-specific user interface can provide more help to the user in the task of managing rules than a general purpose tool such as RMA. If you plan to write your own user interface, refer to the RuleImplementor ([../javadoc/ee/com/ibm/websphere/brb/RuleImplementor.html](http://../javadoc/ee/com/ibm/websphere/brb/RuleImplementor.html)) of the BRBeans Javadoc for assistance.

## Rule management APIs

Business Rule Beans (BRBeans) provide a set of APIs to perform rule management tasks programmatically. These tasks include creating, deleting, and updating rules and folders. These APIs are provided to simplify the interaction with the BRBeans enterprise beans. Use these APIs to perform rule management tasks instead of coding directly in the EJB interfaces.

The rule management APIs consist of the classes in the `com.ibm.websphere.brb.mgmt` package. You might use the following main classes:

### IRule

This is the interface used to access the object representing a business rule in BRBeans. It provides methods to read and update attributes of the rule, to delete the rule, and to make a copy of the rule. The methods to create rules are in the `IRuleFolder` interface since you must always create a rule and specify a particular folder in which it will reside. In the rule management APIs, any time you get a rule, you have the option to receive a reference to the enterprise bean itself or to receive a local copy of the data contained in the enterprise beans. Regardless of which option you choose, the `IRule` interface can be used to access the returned object. If a local copy is requested, it is possible to cast the returned object to an `IRuleCopy`. `IRuleCopy` extends `IRule` and adds a couple additional methods to those defined by `IRule`. See the **IRuleCopy** section for more details.

### IRuleCopy

This is the interface used to access a local copy of the enterprise bean that represents a business rule. An object implementing this interface is returned from rule management API methods if you ask for a local copy of the rule. The main reason for requesting a local copy is performance. Calling a method on a local copy is much faster than calling the method on the actual enterprise bean. If you need to access several different rule attributes, this may make a big difference. Similarly, when updating a rule, all updates can be sent to the enterprise bean in one method call instead of many. The individual set methods are called on the copy and then the `updatePersistentRule()` method is called to actually send the updates to the enterprise bean.

### IRuleFolder

This is the interface used to access the object representing a rule folder. It provides methods to create, delete, and find rules and subfolders. It also provides methods to move and rename the folder, and to get the parent folder. The `IRuleFolder` representing the root folder is generally what you start with when performing rule

management tasks. Once you have the root folder you can navigate up and down the folder hierarchy and access rules contained within the folders.

### **RuleMgmtHelper**

This is a helper class intended to contain methods that are of general use for performing rule management tasks. Currently, the only methods available are used to get the `IRuleFolder` representing the root folder. The root folder is normally the starting point for performing rule management tasks.

### **IParameter**

This is the interface used to represent an initialization or firing parameter stored in a Rule EJB. Every parameter has a user description and a value that are accessible from this interface. The following classes are provided to implement the `IParameter` interface:

#### **ConstantParameter**

This is the most common type of parameter. It represents a single constant value that is to be passed as an initialization or firing parameter.

#### **MethodCallParameter**

This class represents a parameter whose value is determined by calling a method on the target object. The method to call must be a public method and must take zero parameters. This is only used for firing parameters.

#### **TriggerPointParameter**

This class represents a parameter that is retrieved from one of the trigger point firing parameters. This is mainly used for reordering the firing parameters passed on the trigger point. This is only used for firing parameters.

For more details on the rule management interfaces, including a number of coding examples, refer to the `com.ibm.websphere.brb.mgmt` (`../javadoc/ee/com/ibm/websphere/brb/mgmt/package-summary.html`) package in the BRBeans Javadoc.

## **BRBeans performance enhancements**

Externalizing business logic using Business Rule Beans (BRBeans) has many benefits, but does not come completely without a cost. Since every business rule is represented by an enterprise bean, then, in general, every rule trigger is performed in three parts:

1. a query is performed to find the enterprise beans that represent the rules to be triggered.
2. a remote method call is performed on the EJB instance to trigger the rule.
3. a remote method call is made to determine whether to fire the rule locally or remotely.

The first two steps both require server processing so processing can become rather slow.

This section documents the following ways to improve performance:

- Performance enhancements through caching
- Performance enhancements using indexes
- Performance enhancements by changing the firing location

## Performance enhancements through caching

The Business Rule Beans (BRBeans) framework incorporates a cache on the client side; that is, wherever the `trigger()` method on the `TriggerPoint` object is called. This cache is scoped to the Java virtual machine (JVM) in which the client is running so that any trigger calls performed in a particular JVM use the same cache and two triggers performed in different JVMs use two different caches. The BRBeans cache caches the results of all of the queries performed to find a set of rules to be triggered. The next time a trigger is performed in that JVM with the same rules specified, the rules are found in the cache and the query does not require server processing.

Once the rules are found in the cache they are triggered, either locally or remotely, depending on how they were defined. If a rule found in the cache is specified to be triggered locally, then the entire trigger process for that rule is performed on the client without calling the server. Even if the rule is specified to be triggered remotely, finding the rule in the cache eliminates one call to the server since the query is not performed on the server.

The BRBeans cache can improve performance greatly, however it has one disadvantage: changes made to rules are not recognized immediately.

When a change is made to a rule on the server, there is no way to inform all of the potential clients that something has changed and that they may need to refresh their caches. Thus, the client cache must check periodically to see if anything in the persistent rule data has changed. This is implemented by associating a polling frequency with the cache. This polling frequency specifies an interval of time that the cache waits before checking to see if anything has changed. The next time a trigger is performed after a polling interval has passed, the cache checks to see if any changes have been made to the persistent rule data stored on the server. If no changes have been made, then the cache is not refreshed. If any changes have been made, the entire cache is cleared so that the changes are picked up. Thus, changes to the rules are only picked up by the cache after a polling interval has passed.

The default polling frequency is 10 minutes. The user can change this value by changing the single initialization parameter specified for the special rule named `com/ibm/websphere/brb/BRB CacheRule`. The value for this initialization parameter is in the following format: `hh:mm:ss`

`hh` stands for hours, `mm` stands for minutes, and `ss` stands for seconds.

Thus the default of 10 minutes is specified by a value of `00:10:00`. To specify a polling frequency of, for example, 1 hour, 30 minutes, specify `01:30:00`

When this value is changed, it does not take effect until the previous polling interval has passed. Thus, if the previous polling interval is set to 24 hours and the polling frequency is changed to 1 hour, the new frequency does not take effect until the previous 24 hour polling interval passes. The only other ways to get the new frequency to take effect are to either restart the client (since this causes the cache to be re-initialized from scratch) or have the client code call the `refreshCache()` method on the `TriggerPoint` object. If there is more than one client JVM performing triggers, this must be done for each client since each JVM has its own cache.

**Note:** There is only one `BRB CacheRule` and this rule applies to all clients. There is no way to set different polling frequencies for different clients.

Caching can be disabled for a particular TriggerPoint object using the `disableCaching()` method. After the `disableCaching()` method is called any triggers performed using that TriggerPoint object can not use the cache. Triggers performed using other TriggerPoint objects are not affected.

### **Performance enhancements using indexes**

Creating an index over the database table that is used to store rules is an important way to improve the performance of rule queries. It is recommended that an index be created over the `rulename` column of the table containing the rules. This greatly improves the performance of rule-triggered queries that are looking for a rule or rules with a specific name. The index saves the query from searching every row in the table. Refer to the documentation for your database for instructions on how to create an index.

### **Performance enhancements by changing the firing location**

The Business Rule Beans (BRBeans) framework allows you to specify where to fire a particular rule. This determines where the rule implementor is instantiated and invoked. The following lists the possible values for the firing location:

**Local** Fires the Java rule implementor in the same JVM in which the trigger was performed.

#### **Remote**

Fires the Java rule implementor on the server where the rules exist.

#### **Anywhere**

Tries to fire the Java rule implementor locally. If the Java rule implementor cannot be found, then it fires the the implementor remotely. This is the default value.

For simple rule implementors that do not perform any server-intensive work, specifying **Local** usually results in the best performance. This is true both without and with caching. A complete comparison of local firing versus remote firing must consider four cases: local and remote firing without caching and local and remote firing with caching. A description of these four cases follow:

#### **Remote call without caching**

Without caching, the work done to fire a rule remotely involves the following:

1. Finding the rule.
2. Determining whether the rule is to be fired locally or remotely.
3. Calling fire on the remote rule.

Each of these three operations requires a remote call to the server.

#### **Local call without caching**

Without caching the work done to fire a rule locally involves the following:

1. Finding the rule.
2. Determining whether the rule is to be fired locally or remotely.
3. Calling fire on a local copy of the rule.

This requires only two remote calls. Firing locally results in a savings of one remote method call.

#### **Local call with caching**

With caching, local firing results in even more dramatic improvements. The work done to fire a rule remotely involves the following:

1. Finding the rule. This involves a search of the local cache and does not involve calling the database.
2. Calling fire on the remote copy of the rule found in the server's cache.

This requires only one remote method call.

### Remote call with caching

The work done to fire a rule locally with caching involves the following:

1. Finding the rule. This involves a search of the local cache and does not involve calling the database.
2. Calling fire on the local copy of the rule found in the cache.

This does not involve any remote method calls. The entire rule firing process takes place locally without remote method calls. To get the full benefit of the BRBeans cache, use local firing. However, remotely fired rules still benefit from the cache due to the elimination of the query on the server and the elimination of the remote call to determine whether the rule is being fired locally or remotely.

There may be some cases where a rule implementor must perform some work that requires significant interaction with the server. In these cases, it may be beneficial to have rules using this rule implementor defined to be fired remotely. This might make the server interaction performed by the implementor more efficient.

**Note:** In addition to performance, maintenance also must be considered in relation to specifying a firing location. The rule implementor classes for rules that are defined to be fired locally must be present on any client system that tries to fire those rules. Otherwise, the implementor cannot be instantiated when the rule is fired. This can result in maintenance problems when the rule implementors are changed since they must be updated on many different systems.

---

## Developing BRBeans

Although the development tasks in this article are shown in sequential order, the exact sequence is left to your discretion. In particular, you might choose to create the business rules before or after the trigger point is placed. Also, you can write the rule implementor before or after creating the actual business rules. However, if you do not have the rule implementor at the time that the business rule is created, then you cannot complete the rule implementor field or the initialization parameters in the rule. However, you can leave the business rule marked as not ready until you can complete that information. For this reason, we have chosen to list the task of writing the rule implementor first. Complete the following tasks to develop Business Rule Beans:

### Steps for this task

1. Determining where to place the trigger point  
Determine where there are points of variability and where business logic must be externalized as part of application analysis and design process.
2. Placing a trigger point in the application code  
Add code to the application to invoke the trigger point framework, to find business rules, and to fire the rules.
3. Administering strategy objects to control triggers

Control the process of triggering the rules using a set of strategy objects.

#### 4. Implementing the business rules

Use a rule implementor, written in Java, to implement the BRBeans RuleImplementor interface. Also, create business rules invoked by the rule implementor.

#### What to do next

It is possible to develop your own customized strategy objects and rule implementors. See Customized strategy objects and Customized rule implementors for more information.

## Determining where to place a trigger point

To determine where to place your trigger points, you can use either the case analysis or the Object Interaction Diagrams (OIDs) method. The following are examples of methods that you can use to determine where to place a trigger point:

### Use case analysis

Trigger points can be found during analysis by inspecting the use cases or user interaction scenarios that are typically developed as statements of requirement as input to the analysis process. A fragment of a use case is shown below:

*The vehicle is entered into the system or chosen. The customer service representative attempts to locate the named driver in the system. If the driver is not found, she or he is added to the system and then picked.*

*Otherwise the found driver is simply picked. If the vehicle is an auto, anyone between the ages of 16 and 75 can be picked as a driver. If the vehicle is a truck, only drivers 16 to 70 years old can be picked. And if the vehicle is a motorcycle, drivers 14 to 65 can be picked.*

*After the driver has been picked, a rate quote can be performed...*

To identify potential trigger locations in use case analyses such as this one, look for **keywords** such as:

- "if X is in a special category Y" (For example, "if the vehicle is a truck")
- "except when"
- "unless"
- "depends on"

### Object Interaction Diagrams (OIDs)

OIDs that are based on use cases can yield a number of observable patterns that can be used to identify trigger points fairly easy. The following are some of the rules to look for and where the trigger point might be placed:

- Validation of edits on create methods.
- Validation of edits on set methods.
- Referential integrity of edits on methods that set references.
- Cardinality checks at a consistency point (a point in time where all of the data is expected to be self consistent).
- Required fields checks at a consistency point.
- Cross field edits at a consistency point.

- Constraints or derivations that have a high potential for reuse (especially if the algorithm is complex) at any appropriate point.
- Constraints or derivations that a business desires to be consistent across applications (at any appropriate point).
- Constraints or derivations where the business wants to decouple the maintenance cycle for a rule from the maintenance cycle for the code (at any appropriate point).

### Results

By using either method, you will be able to identify where to locate trigger points to use Business Rule Beans (BRBeans) effectively.

### What to do next

Once the trigger points have been identified, place the trigger point or points. See [Placing a trigger point](#) for more information.

## Placing a trigger point in the application code

### Before you begin

Before placing a trigger point, review the following topics:

- Trigger point framework
- Determining where to place a trigger point

The `TriggerPoint` class is the primary interface of the Business Rule Beans (BRBeans) Trigger Point framework. The class is used to transfer control to the Trigger Point framework to find and fire the rules specified in the application's trigger point.

### Steps for this task

1. Create an instance of the `com.ibm.websphere.brb.TriggerPoint` class. All rule triggers must be performed against an instance of this class. Also, set any desired strategies on the `TriggerPoint` instance.
2. Gather together the parameters to be passed on the trigger.

For the simple `trigger()` and `triggerClassifier()` methods this includes the following:

#### **An optional target object**

This can be used to specify an object that is to be the target of the rule's algorithm. Whether this is needed depends completely on the design of the rule implementor being used.

#### **The firing parameters for this rule trigger**

This is an array of run-time parameters needed by the rule to satisfy its business purpose. The exact set of required firing parameters is determined by the rule implementor that is used by the rule.

**Note:** Any firing parameters defined on the rule itself will override whatever is passed here.

#### **Information identifying the rule or rules to be triggered**

Normally this is either a single `String` containing the name of the rule to be triggered or an array of `Strings` each element of which is the

name of a rule to be triggered. However if a custom finding strategy is being used, this could be whatever information it needs in order to find the correct rules.

The `triggerSituational` method differs in that it takes two sets of firing parameters and two sets of rule identification information: one set for the classifier rules and one for the classified rules.

3. Invoke the `trigger()`, `triggerClassifier()`, or `triggerSituational()` method of the `TriggerPoint` instance. This will actually trigger the rule or rules.
4. Process the results of the triggered rule or rules.

#### What to do next

Examples of how to code a trigger point call are provided in the following topics:

- Simple trigger point
- Classifier trigger point
- Situational trigger point

For a detailed description of the trigger point programming interfaces refer to the Trigger Point class ([../javadoc/ee/com/ibm/websphere/brb/TriggerPoint.html](http://../javadoc/ee/com/ibm/websphere/brb/TriggerPoint.html)) in the BRBeans Javadoc.

## Administering strategy objects to control triggers

Strategy objects are used to alter `TriggerPoint` functions. The two simple trigger methods, `trigger()` and `triggerClassifier()`, perform their function in the following sequence:

#### Steps for this task

1. Find the rules.
2. Filter out those rules which are not desired.
3. Fire the remaining rules.
4. Combine the results and return to the caller.

#### Results

The complex trigger method `triggerSituational()` does this sequence of steps twice; the first sequence finds the classification to feed into the second sequence.

Default strategy objects already are defined for each of the four `TriggerPoint` steps and they are used if none are specified explicitly. For each of these steps, there are at least two strategy objects used, one for triggering classifier rules, and one for triggering non-classifier rules. For the filtering step, there are actually three pairs of strategies that are used, based on the number of rules which the finding strategy returns (zero, one, or multiple).

While the sheer number of strategies that are available can be intimidating (twelve different strategy classes can be set), very few will need updating. Most users will modify the filtering strategies or the combining strategies only.

A number of predefined strategy objects are provided and can be used for the majority of cases. Although the following strategies are described separately, they have a definite dependency on each other. For example, the `FilteringStrategy` filters



rules from the FindingStrategy; the FiringStrategy uses the results of the FindingStrategy to operate; and the CombiningStrategy depends on the results of the FiringStrategy.

#### **FindingStrategy**

The FindingStrategy accesses the data store and returns those rules that meet the search criteria specified. See the FindingStrategy method for more information.

#### **FilteringStrategy**

The FilteringStrategy takes the list of rules that were found by the FindingStrategy and filters out the rules that should not be fired. See the FilteringStrategy method for more information.

#### **FiringStrategy**

The FiringStrategy takes the rules that were found by the FindingStrategy, (possibly modified by the FilteringStrategy), fires them each in order, and returns an array containing the results of each rule. See the FiringStrategy method for more information.

#### **CombiningStrategy**

The CombiningStrategy takes the results of the rules that are fired by the FiringStrategy and combines them to form a reasonable result to the TriggerPoint caller. See the CombiningStrategy method for more information.

The Java classes for these strategy objects are defined in the `com.ibm.websphere.brb.strategy` package. Static constants also are defined in the interfaces for the various strategies. This allows easy access to instances of the strategy classes to set them on the TriggerPoint.

Also, it is possible to write your own strategy class if the supplied ones do not perform the function you need. See Customized strategy objects for more details.

## **Implementing business rules**

### Before you begin

After determining where to place a trigger point, placing the trigger point, and defining your strategy objects, you must provide a method to implement the business rules and then create the rules. In addition, you might choose to organize your rules by creating rule folders.

**Note:** Although the development tasks in this article are shown in sequential order, the exact sequence is left to your discretion. In particular, you might choose to write the rule implementor before or after creating the actual business rules.

The business rule encapsulates and externalizes the business logic for the rule and any data that parameterizes the rule. Complete the following process to implement business rules:

### Steps for this task

1. Provide an implementation, called a rule implementor, for each business rule that you create

The rule implementor provides the actual business logic for the rule, implemented in Java. The rule implementor's `fire()` method is called when the business rule is triggered to actually perform the processing for the rule. Several generic rule implementors are shipped with BRBeans, which might be

useful in some situations. However, if these rule implementors do not meet your needs, you can write your own rule implementor. Refer to the section on writing your own rule implementor for details.

2. Use the Rule Management Application (RMA) to create the business rule
  - a. In the **Rule Browser** window, select the folder where you want the new rule to be created.
  - b. From the main menu, click **File > New > Rule**.  
In the **New Rule** properties window, use the following tabs to define the rule:
    - General**  
Use this tab to enter general information about the rule.
    - Implementation**  
Use this tab to define the manner in which the rule is implemented.
    - Description**  
Use this tab to define the purpose and intent of the rule.
    - Dependent Rules**  
Use this tab to specify the rules that the newly created rule will depend upon.
    - Other** Use this tab to to establish precedence, and enter information that is relevant to you, but does not fit into any other category.
  - c. To complete the creation of the rule, click **OK**.  
If there are any mandatory fields still undefined, either go back and give them a value, or make the rule unavailable for use (see **Status** in the **General** tab for more information on this).
3. **(Optional)** Create a rule folder using the Rule Management Application to contain the new business rule
  - a. In the **Rule Browser** window, select the folder where you want the new folder to be nested.
  - b. From the main menu, click **File > New > Folder**.  
A new folder appears in the folder hierarchy in edit mode. Enter a folder name and hit the **Enter** key.

---

## Assembling applications for use with BRBeans

When you are ready to ship your application, include a Business Rule Beans (BRBeans) JAR file in your EAR file. There are several of these JAR files in the <WAS\_HOME>/BRBeans directory; one for each supported database. Each name reflects the database type that it uses (for example, BRBeansDB2.jar). These JAR files contain three enterprise beans with the following JNDI names:

- brbeans/application/Rule
- brbeans/application/RuleFolder
- brbeans/application/RuleHelper

In your EAR files, complete the following steps:

### Steps for this task

1. Change the JNDI names of the BRBeans enterprise beans to make them unique for your application.

For example, if your application is called MyApp, you could change the first one to brbeans/MyApp/Rule or com/MyCompany/MyApp/Rule.

2. Define EJB references to these three enterprise beans in any module where a `trigger...()` method exists in one or more of its classes.

You can do this using the WebSphere Studio Application Developer tool or the Application Assembly tool. The Name field should contain the following and correspond to the enterprise beans listed above:

- `ejb/com/ibm/ws/brb/Rule`
- `ejb/com/ibm/ws/brb/RuleFolder`
- `ejb/com/ibm/ws/brb/RuleHelper`

**Note:** The JNDI name on the **Bindings** tab should be the same as the JNDI names that you gave earlier to the enterprise beans.

Since the BRBeans enterprise beans refer to each other, there are also EJB references defined within the BRBeans JAR file itself. Each of the three BRBeans enterprise beans has two EJB references defined. These six references also need to be updated with the JNDI names you assigned earlier.

3. If you are not using the `BRBeansCloudscape.jar` file, skip this step. In the `BRBeansCloudscape.jar` file each entity enterprise bean has a resource reference defined for the data source it will use. You must update the JNDI binding for this reference to specify the JNDI name of the actual data source you want to use for the BRBeans entity enterprise bean. The shipped `BRBeansCloudscape.jar` file contains a dummy JNDI name for each binding. There are two resource references that need to be updated: one for the Rule enterprise bean and one for the RuleFolder enterprise bean. You can use the WebSphere Studio Application Developer tool or the Application Assembly tool to perform this update.

### What to do next

After you have an EAR file that includes your application code and the BRBeans enterprise beans, complete the following steps to install and run the application:

1. Deploy the EAR file to generate run-time code for the BRBeans enterprise beans and any other enterprise beans that you may have in your application. Deployment can be accomplished either as a separate step (using the `ejbdeploy` command) or as part of the process of installing the EAR file onto an application server. Refer to the base WebSphere documentation for more information about deploying an EAR file.
2. The deployment process generates a file named `Table.ddl`. `Table.ddl` contains database commands to create the database tables needed by the BRBeans entity enterprise beans. You must use your database software to create a database and then use the commands in the `Table.ddl` file to create tables in this database. Refer to your database software documentation for more information on the commands needed to set up a database and the database tables.
3. Create a JDBC provider and a data source to access the database created in the previous step. For all of the databases except IBM Cloudscape, the BRBeans enterprise beans are configured to require a Version 4 data source. For Cloudscape, you can use a normal data source.

**Note:** The JNDI name you specify for the data source also must be specified when you install the EAR file on the application server so that the server knows which data source to use. Refer to the base WebSphere documentation for more information on these topics.

4. Install the EAR file on an application server. To install the EAR file, either use command line tools or the WebSphere Administrative Console.

**Note:** The EAR file can be deployed either as a separate step or as part of the installation process on the server. Refer to the base WebSphere documentation for more information on installing an EAR file on an application server.

5. Start the application using the WebSphere Administrative Console.

You now have an application installed and running using BRBeans.

---

## Managing rules

### Before you begin

In Business Rule Beans (BRBeans), rule management involves making changes to the set of business rules used by applications. This can include any of the following activities:

- Implementing the business rules: Creating rules and rule folders
- Copying and moving rule or rule folders
- Working with Quick Copy
- Finding a rule
- Deleting rules
- Deleting rule folders
- Changing the properties of a rule
- Importing a rule
- Exporting a rule
- Renaming rules
- Renaming rule folders
- Formatting columns
- Changing the date and time format

There are two different interfaces that can be used for rule management:

### **Rule Management Application**

An external user interface that allows users to manage rules interactively. It provides a general purpose interface for managing rules where no assumptions are made about the content or implementation of the rules. For information on how to use the Rule Management Application, see *Starting the Rule Management Application*.

### **Rule management APIs**

A programmatic interface that can be used by programmers to manage rules or to customize an external user interface. For more information on Rule management APIs, see the `com.ibm.websphere.brb.mgmt` (`../javadoc/ee/com/ibm/websphere/brb/mgmt/package-summary.html`) package.

Rules can be managed in any way that makes sense for your application, but the BRBeans framework was designed with the following administrative paradigm in mind:

### Steps for this task

1. Understand the desired change in business behavior.
2. Inspect the application documentation (in particular information that indicates where trigger points are located) to understand where the changes need to be made in the system.

3. Inspect the corresponding set of existing business rules using the Rule Management Application (or your own custom management application, if you have one) to understand which rules need to change.
4. Use the Rule Management Application, on a test system, to create one or more new rules that implement the required new behavior. Give these rules the correct name so that they are triggered by the appropriate trigger point. Also, make sure that these new rules are currently in effect.
5. Withdraw (by setting the end date of the rule), on the test system, all of the rules that are to be superseded.
6. Test the application to ensure that it behaves as expected.
7. Export the new rules using the Rule Exporter on the test system. Schedule the rules to become effective at the correct point in time.
8. Export the rules to be superseded using the Rule Exporter on the test system. Set them to expire when the new rules come into effect.
9. Import the new rules using the Rule Importer on the production system. This creates the new rules and schedules them to become effective at the date you specified when you exported them.
10. Import the rules to be superseded using the Rule Importer on the production system. This puts the new end date into the existing rules on the production system and sets them to expire on the specified date.

#### What to do next

For more information on the overall topic of Externalizing business rules, see [Using Business Rule Beans](#)

## Starting the Rule Management Application

### Before you begin

Review the overview topic [Managing rules](#)

To administer BRBeans, use the Rule Management Application (RMA). To launch the RMA, complete the following steps:

### Steps for this task

1. Open a command window and change to the following directory:  
`<WAS_HOME>/bin`
2. For Microsoft Windows platforms, enter `rulemgmt.bat <properties-file>`. For UNIX platforms, enter `rulemgmt.sh <properties-file>`  
where `<properties-file>` is a fully qualified name of a file containing port, host, and JNDI names used for the Business Rule Beans (BRBeans) enterprise beans. If you are using localhost port 2809 and you are using the default JNDI names for the BRBeans enterprise beans, you can specify the following in the default properties file: `<WAS_HOME>/bin/brbeansDefaultProperties`. For a full definition of the contents of this file, see the BRBeans properties file.

## Copying or moving rules or rule folders

Copy or move rules or rule folders either by cutting and pasting, or dragging and dropping.

To copy or move rules or rule folders, proceed as follows:

### Cutting and pasting

Use the menu commands (**Edit > Copy**, **Edit > Cut** and **Edit > Paste**) or keyboard commands (**CTRL-C**, **CTRL-V**, and **CTRL-X**).

### Dragging and dropping

Highlight the rule or rule folder you want to copy. Then press and hold the right mouse button, drag the cursor to the target location, and release. Select **Copy** or **Move** from the list.

**Note:** A rule also can be copied so that the copy replaces the existing rule at a specified date. This is referred to as a Quick Copy.

See Managing rules for other tasks related to the management of your rules.

## Working with Quick Copy

Use **Quick Copy** to make a copy of a rule that will replace the existing one on a specified date.

For example, suppose that we have an "isSeniorCitizen" rule. Currently a person is considered a senior citizen if they are 62 years of age or older. Starting on January 1, 2002, we are going to change this to 65. Use Quick Copy to specify the new date, and change the age from 62 to 65. The current rule will be set to expire on the same date that the new rule will take effect with the new senior citizen age defined as 65.

To use **Quick Copy** from the Rule Browser or Search Results window, proceed as follows:

#### Steps for this task

1. Select the rules you want to **Quick Copy**.
2. From the main menu, click **Edit > Quick Copy**.
3. In the Quick Copy window, specify in the following fields how the copy will differ from the original:

#### **Start Date For New Rule**

Enter the date that the new rule will replace the existing rule. Use the date and time format that is shown.

#### **Change parameter values for new rule**

Enter the new parameter values.

4. Click **OK** to finish.

#### What to do next

See Managing rules for other tasks related to the management of your rules.

## Finding a rule

To search for a specific rule using the Rule Management Application **Find** function, proceed as follows:

#### Steps for this task

1. To search through all rules in all folders:
  - a. From the main menu of the Rule Browser, click **Edit > Find**.
  - b. Use the tabs in the Find Rules window to determine your search criteria.

2. To search a specific folder:
  - a. Right-click the folder and select **Find** from the list.
  - b. Use the tabs in the Find Rules window to determine your search criteria.

#### What to do next

The results of your search are displayed in a Search Results window.

See Managing rules for other tasks related to the management of your rules.

## Deleting rules

To delete rules from the Rule Browser or Search Results window, proceed as follows:

#### Steps for this task

1. Select the rules you want to delete.
2. From the main menu, click **File > Delete**.
3. Click **Delete** and then confirm the delete request.

#### What to do next

**Note:** You cannot delete `com/ibm/websphere/brb/BRB CacheRule` as this rule is needed by the Business Rule Beans run-time environment.

See Managing rules for other tasks related to the management of your rules.

## Deleting rule folders

To delete rule folders from the Rule Browser window, proceed as follows:

#### Steps for this task

1. Select the folder you want to delete.
2. From the main menu, click **File > Delete**.
3. Click **Delete** and then confirm the delete request.

#### What to do next

**Note:** You cannot delete the root folder or any of the folders in the path `com/ibm/websphere/brb`.

See Managing rules for other tasks related to the management of your rules.

## Changing the properties of a rule

To change the properties of a rule, perform the following steps in either the Rule Browser or Search Result window:

#### Steps for this task

1. Highlight the rule you wish to edit.
2. From the main menu, click **File > Properties**.

In the Rule Properties properties window, use the following tabs to edit the rule's definition:

**General**

Use this tab to edit general information about the rule.

**Implementation**

Use this tab to edit the manner in which the rule is implemented.

**Description**

Use this tab to edit the purpose and intent of the rule.

**Dependent Rules**

Use this tab to edit the list of rules that the newly created rule will depend upon.

**Other** Use this tab to establish precedence, and enter information that is relevant to you, but does not fit into any other category.

3. To complete the editing of the rule, click **OK**.

If there are any mandatory fields still undefined, either go back and give them a value, or make the rule unavailable for use (see **Status** in the **General** tab for more information on this).

What to do next

See Managing rules for other tasks related to the management of your rules.

## Importing a rule

To import rules from an XML format, use the Rule Browser window and proceed as follows:

Steps for this task

1. In the main menu, click **File > Import**.
2. In the Import Rules window, specify the file you want to import.
3. Click **OK**.

Rules and rule folders are created as specified within the XML.

What to do next

See Managing rules for other tasks related to the management of your rules.

## Exporting a rule

To export rules, use the Rule Browser or Search Results window and proceed as follows:

Steps for this task

1. In the main menu click **File > Export**.
2. In the **Export Rules Wizard**, proceed as follows:
  - a. In the Specify Rules to Export window, select the rule or rules that you want to export and click **Next**.
  - b. In the Change Effective Dates On Exported Rules window, alter the start and end dates of the rule (if desired) and click **Next**.
  - c. In the Select File For Rule Export window, choose a name and location for the exported rule.
3. Click **Export** to finish.

What to do next



See Managing rules for other tasks related to the management of your rules.

## Renaming rules

To rename rules, use the Rule Browser or Search Results window and proceed as follows:

### Steps for this task

1. Highlight the rule you want to rename.
2. From the main menu, click **File > Rename**.
3. Type a new name and press the **Enter** key.  
To cancel the name change while it is still in progress, press the **Esc** key.

### What to do next

See Managing rules for other tasks related to the management of your rules.

## Renaming rule folders

To rename rule folders, use the Rule Browser or Search Results window and proceed as follows:

### Steps for this task

1. Place the folder name in edit mode by performing one of the following tasks:
  - a. Right-click the folder and select **Rename** from the list.
  - b. Highlight the folder and click **File >Rename** in the main menu.
  - c. Click the folder name twice with a slight pause between each click.
2. Type a new name and press the **Enter** key.  
To cancel the name change while it is still in progress, press the **Esc** key.

### Results

**Note:** You cannot change the name of the root folder.

### What to do next

See Managing rules for other tasks related to the management of your rules.

## Specifying columns

To choose which columns you want shown in your Rule Browser window, perform the following steps in either the Rule Browser or Search Results window:

### Steps for this task

1. From the main menu, click **View > Specify Columns**.
2. In the Specify Column window, proceed as follows:

#### **To add a new column**

Select one or more entries in the **Available columns** list and click the **Add** button. The selected entries are added to the end of the Columns displayed list.

#### To remove a column

Select one or more entries in the Columns displayed list and click the **Remove** button. The selected entries are added to the end of the Available columns list.

#### To reorder columns

Select one or more entries in the Columns displayed list. To move the entries towards the top of the list, click the Up arrow. To move them towards the bottom, click the Down arrow.

#### What to do next

See Managing rules for other tasks related to the management of your rules.

## Changing the date and time format

To change the date and time format, use either the **Rule Browser** or **Search Results** window and proceed as follows:

#### Steps for this task

1. In the main menu, click **View > Specify Date/Time Format**.
2. In the **Specify Date/Time Format** window, choose one of the following radio button options:

##### **Use default format for this locale**

Use this option to adjust the date and time format to match the default setting of your current locale.

##### **Select a predefined format for the date and time**

Use this option to choose from one of several existing date and time formats.

##### **Specify a custom format for the date and time**

Use this option to determine your own format for your date and time display.

3. When the example in the lower left of the window meets your needs, click **OK** to finish.

#### What to do next

See Managing rules for other tasks related to the management of your rules.

---

## Rule Browser

The Rule Browser is the main window of the Rule Management Application (RMA), which is the tool used to administer business rules for Business Rules Beans. The RMA is a simple graphic user interface that assists the user in the high-level administration of rules and rule folders. The column on the left side of the Rule Browser window shows a nested hierarchy of all of the existing rule folders. Open one of these rule folders to display the rules. The names of these folders appear in the right column.

To navigate through the information, perform the following actions:

- Click the + icon to expand the folder by one level; click the - icon to collapse it.
- Click a file name to highlight it, right-click it to launch a list of actions, or select an option from the main menu.

## File menu

This article describes the options available in the File menu window.

### New

Use the **New** option to create one of the following objects:

**Folder** The **Folder** selection creates a new folder within the folder currently selected in the browser. To create a new folder, complete the following steps:

1. Select the folder, in the Rule Browser window in which to nest the new folder.
2. Click **File > New > Folder** from the main menu. A new folder appears in the folder hierarchy in edit mode.
3. Enter a folder name and press **Enter**.

**Rule** The **Rule** selection creates a new rule within the folder that is currently selected in the browser. To create a new rule, complete the following steps:

1. Select the folder, in the Rule Browser window, where you want to create the new rule.
2. Click **File > New > Rule** from the main menu.
3. Use the following options, in the New Rule Properties window, to define the rule. For more information on each of these tabs, see the associated help file.

#### General

Use the **General** tab to enter general information about the rule. For more information, click the **New Rule properties window: General tab** link under Related reference.

#### Implementation

Use the **Implementation** tab to define the manner in which the rule is implemented. For more information, click the **New Rule properties window: Implementation tab** link under Related reference.

#### Description

Use the **Description** tab to define the purpose and intent of the rule. For more information, click the **New Rule properties window: Description tab** link under Related reference.

#### Dependent Rules

Use the **Dependent Rules** tab to specify the rules that the newly created rule will depend upon. For more information, click the **New Rule properties window: Dependent Rules tab** link under Related reference.

**Other** Use the **Other** tab to establish precedence and enter information that is relevant to you, but does not fit into any other category. For more information, click the **New Rule properties window: Other tab** link under Related reference.

**Note:** When you change the properties of a rule and there are undefined mandatory fields, either give them a value or make the rule unavailable for use. See "status" on the General tab for more information.

4. Click **OK** to complete the creation of the rule.

## Rule Browser Window

This selection opens a new Rule Browser window on your desktop and shows the content of the currently selected folder.

## Import

Use the **Import** option to import rules that are defined in a file and written in XML. The rules are imported into folders as specified within the XML.

### Importing rules

You can use the Rule Browser window to import from an XML format. To import rules, complete the following steps:

1. Click **File > Import** in the main menu.
2. Specify the file, in the Import Rules window, that you want to import.
3. Click **OK**. The rules are imported as specified within the XML.

For more information, see the **Import Rules window** link under Related reference.

## Export

Use the Export option to export a file in an XML format.

### Exporting rules

You can export rules from the Rule Browser or Search Results windows. To export rules, complete the following steps:

1. Click **File > Export**, in the main menu.
2. Proceed as follows using the Export Rules Wizard:
  - In the Select Rules to Export window, select the rule or rules that you want to export and click **Next**.
  - In the Change Effective Dates On Exported Rules window, alter the start and end dates of the rule, if desired, and click **Next**.
  - In the Select File For Rule Export window, choose a name and location for the exported rule.
3. Click **Export** to finish.

For more information on this Export Rules Wizard, see the **Select Rules to Export window** link under Related reference.

## Delete

Use the **Delete** option to delete the selected rule or rule folder. If a rule folder is selected, all of the rules and subfolders the folder contains also are deleted.

### Deleting rules

You can delete rules from the Rule Browser or Search Results windows. To delete rules, complete the following steps:

1. Select the rules that you want to delete using the Rule Browser window.
2. Click **File > Delete** from the main menu.
3. Click **Delete** and then confirm the delete request.

**Note:** You cannot delete `com/ibm/websphere/brb/BRB CacheRule` as this rule is needed by the Business Rule Beans run-time environment.

### Deleting a folder

To delete a folder, complete the following steps:

1. Select the folder you want to delete using the Rule Browser window.
2. Click **File > Delete** in the main menu.
3. Click **Delete** and then confirm the delete request.

**Note:** You cannot delete the root folder.

## Rename

Use the **Rename** option to rename the selected rule or rule folder.

### Renaming a rule

You can rename rules from the Rule Browser or Search Results windows. To rename a folder, complete the following steps:

1. Highlight the rule you want to rename.
2. Click **File >Rename** from the main menu.
3. Type a new name and press **Enter**. To cancel the name change while it is still in progress, press **Esc**.

### Renaming a folder

1. Place the folder name in edit mode by doing any of the following:
  - Right-click the folder and select **Rename** from the list.
  - Highlight the folder and click **File >Rename** in the main menu.
  - Click the folder name twice with a slight pause between each click.
2. Type a new name and press **Enter**. To cancel the name change while it is still in progress, press **Esc**.

**Note:** You cannot rename the root folder.

## Properties

Use the **Properties** option to modify the properties of the selected rule or rule folder and then click **OK**.

Use the following tabs, in the **Rule Properties** window, to define the rule. For more information on each of these tabs, see the associated help file.

### General

Use the **General** tab to enter general information about the rule. For more information, click the **Rule properties window: General tab** link under Related reference.

### Implementation

Use the **Implementation** tab to define the manner in which the rule is implemented. For more information, click the **Rule properties window: Implementation tab** link under Related reference.

### Description

Use the **Description** tab to define the purpose and intent of the rule. For more information, click the **Rule properties window: Description tab** link under Related reference.

### Dependent Rules

Use the **Dependent Rules** tab to specify the rules that the newly created rule will depend upon. For more information, click the **Rule properties window: Dependent Rules tab** link under Related reference.

**Other** Use the **Other** tab to establish precedence and enter information that is relevant to you, but does not fit into any other category. For more information, click the **Rule properties window: Other tab** link under Related reference.

**Note:** When you change the properties of a rule and there are undefined mandatory fields, either give the fields a value or make the rule unavailable for use. See "status" on the **Rule properties window: General tab** for more information.

### **Validate**

Use the **Validate** option to verify the correctness of the selected rules or folder. A report is generated and displayed in system browser window. Use the browser functionality to view, save, and print the results. The following items are validated:

#### **Start and End dates**

Specifies a valid start date. The end date must be empty (never expires) or be later than the start date.

#### **Java rule implementor name**

Specifies a valid and existing Java rule implementor name. The specified class must implement the RuleImplementor interface.

#### **Dependent rules**

Verifies the existence of each specified dependent rule.

When this option is selected, the Validate Rules window displays. You must specify a Java class path to validate rule implementor classes.

### **Close**

Use the **Close** option to terminate the application and close the window.

### **New Rule properties window: General tab**

Use the **New Rule properties window: General** tab to enter general information about the rule. The following fields and options are available on this tab:

#### **Name and location**

##### **Folder name**

(Required) Use the **Folder name** field to identify the folder in which to create the rule. To browse the existing folders, click the **ellipses** button to the right of the text field.

**Name** (Required) Use the **Name** field to give the rule a name. The name cannot include the forward slash '/' and must contain at least one non-blank character. The name cannot exceed the maximum length of the rule name column in the database table.

#### **Period when rule is in effect**

##### **Start date**

(Required) Use the **Start date** field to specify the date and time that the rule will go into effect. If you do not specify the time, a value of midnight is used.

##### **End date**

Use the **End date** field to determine the date when this rule expires. If you do not specify a value, the rule never expires.

#### **Classification**

Use this section to specify your rule's classification status. Choose one of the following options:

- Rule is not classified and does not perform a classification (default)
- Rule performs a classification
- Rule is classified with the following classification

### Status

#### Rule is available for use

Select **Rule is available for use** when the rule is available for use by the Business Rule Beans run-time environment. This feature is useful when you have not finished creating the rule, but need to save your changes.

### New Rule properties window: Implementation tab

Use the **New Rule properties window: Implementation** tab to define the manner in which the rule is implemented.

**Note:** The rule contains the data, but it does not perform the implementation. Rather, the Java rule implementor implements the rule.

For example, suppose you want to create a rule that determines whether a given person is a senior citizen, 62 years old or older. To implement this rule, give the "com.ibm.websphere.brb.implementor.RuleGreaterThanEqual" Java rule implementor the value 62 and specify it as an initialization parameter. When the application fires the rule, the person's age is passed to the Java rule implementor as a firing parameter and 62 is passed as an initialization parameter. The person's age is compared against the initialization parameter of 62 and a value of true or false is returned from the Java rule implementor to the application. To change the age at which a person is considered a senior citizen, change the value of the initialization parameter.

The following fields and options are available on this tab:

**Java rule implementor:** (Required) Use the **Java rule implementor** field to specify a class to implement this rule. The initialization and firing parameters that are required are determined by looking at the documentation for the specified Java rule implementor.

**Firing location:** Use the **Firing location** field to determine where the rule is fired. You can fire the Java rule implementor on the server where the rules exist or fire it locally on the client machine. The client can be a servlet running on the server. Specify one of the following values for the firing location:

**Local** Use the **Local** option to fire the Java rule implementor local to the application that fired the rule.

#### Remote

Use the **Remote** option to fire the Java rule implementor on the server where the rules exist.

#### Anywhere (default)

Use the **Anywhere** option to attempt to fire the Java rule implementor locally first. If the Java rule implementor cannot be found, then it is fired remotely.

To choose the value of the firing location, you must take both performance and maintenance into consideration. Most rules perform better if they are run on the same Java virtual machine (JVM) as the application (locally). However, there might

be cases where a Java rule implementor performs server-intensive tasks, in which case the rules might run better when they run on the server. To run locally, you must have all of the Java rule implementors installed locally. They must be accessible by the application that fires the rules.

**Initialization parameters:** The initialization parameters contain constant values passed to the rule implementor when it is initialized. Typically, the initialization parameters contain values that might change as your business practices evolve, such as the age at which a person is considered a senior citizen or the current interest rate for a loan.

To add a new initialization parameter to the list, click **Add** and fill in the fields in the Add Initialization Parameter window.

To edit an existing initialization parameter, highlight it, click **Change**, and fill in the fields in the Change Initialization Parameter window.

To delete an initialization parameter, highlight it and click **Delete**.

To change the order of the initialization parameters, highlight one and click the up or down arrows to move it to a new location.

**Firing parameters:** The firing parameters contain values passed from the trigger point in the application to the Java rule implementor at run time. You can alter the parameters coming from the application before passing them to the Java rule implementor. Typically, these parameters are left unaltered.

For example, when implementing the "isSeniorCitizen" rule that determines whether a person is a senior citizen, you might want to pass a person from the application to the rule as the target object. However, the "isSeniorCitizen" rule uses the RuleGreaterThanEqual Java rule implementor, which requires that you pass an integer value. You can alter the firing parameters to specify that the method "getAge" is called on the person object and pass that result to the Java rule implementor.

You must choose one of the following options:

**Pass firing parameters from trigger point unchanged**

The parameters specified in the trigger point of the application are passed to the Java rule implementor without being altered. This is the default value.

**Specify firing parameters**

The values specified in the table are passed to the Java rule implementor.

To add a new firing parameter to the list, click **Add** and fill in the fields in the Add Firing Parameter window.

To edit an existing firing parameter, highlight the firing parameter, click **Change**, and fill in the fields in the Change Firing Parameter window.

To delete a parameter, highlight the firing parameter and click **Delete**.

To change the order of the firing parameters, highlight one and click the up or down arrows to move it to a new location.



## Add Initialization Parameter window

Use the Add Initialization Parameter window to add initialization parameters to a rule. The initialization parameters contain constant values passed to the Java rule implementor when it is initialized. Typically, the initialization parameters contain values that will change as your business practices evolve. These values might be the age at which a person is considered a senior citizen or the current interest rate for a loan.

To add an initialization parameter, proceed as follows:

1. Fill in the fields as needed.
2. Click **Add**. The initialization parameter is added and the window remains open to specify additional parameters.
3. When you are finished specifying initialization parameters, click **Close**.

The following fields and options are available in this window:

**Description:** Use the **Description** field to specify a description of the initialization parameter. This field can contain any information necessary to describe the purpose of the initialization parameter.

**Type:** Use the **Type** field to specify the type of data that is contained within this initialization parameter. The data itself is stored in the **Value** field (see the following field description). For example, if this initialization parameter specifies the age at which a person is considered a senior citizen, then the Type likely is an "Integer". If the initialization parameter specifies a company name, such as "IBM", then the Type is a "String".

The following values are available:

- String
- Character
- Byte
- Short
- Integer
- Long
- Float
- Double
- Boolean
- java.math.BigDecimal
- java.math.BigInteger
- Null Value

**Value:** Use the **Value** field to add a value for the parameter. For example, if the initialization parameter is intended to specify the age at which a person is considered a senior citizen, then this field might be set to 62.

## Change Initialization Parameter window

Use the Change Initialization Parameter window to edit an existing initialization parameter.

The initialization parameters contain constant values passed to the Java rule implementor when it is initialized. Typically, the initialization parameters contain

values that change as your business practices evolve. These values might be the age at which a person is considered a senior citizen or the current interest rate for a loan.

To edit an existing initialization parameter, modify the fields and click **OK**.

The following fields are available in this window:

**Description:** Use the **Description** field to specify a description of the initialization parameter. This field can contain any information necessary to describe the purpose of the initialization parameter.

**Type:** Use the **Type** field to specify the type of data that is contained within this initialization parameter. The data itself is stored in the **Value** field (see the following field description). For example, if this initialization parameter specifies the age at which a person is considered a senior citizen, then the Type likely is an "Integer". If the initialization parameter specifies a company name, such as "IBM", then the Type is a "String".

The following values are available:

- String
- Character
- Byte
- Short
- Integer
- Long
- Float
- Double
- Boolean
- java.math.BigDecimal
- java.math.BigInteger
- Null Value

**Value:** Use the **Value** field to add a value for the parameter. For example, if the initialization parameter is intended to specify the age at which a person is considered a senior citizen, then this field might be set to 62.

### **New Rule properties window: Description tab**

Use the **New Rule properties window: Description** tab to define the purpose and intent of the rule. All of the fields in this panel are optional and none of them are used by the Business Rule Beans run time. The following fields are available on this tab:

**Business Intent:** Use the **Business Intent** field to describe the business intent of this rule.

**Description:** Use the **Description** field to define a general description of the rule and its purpose.

**Original requirement:** Use the **Original requirement** field to compose a description of the original requirement that created this rule.

## New Rule properties window: Dependent Rules tab

Use the **New Rule properties window: Dependent Rules tab** to specify the rules that the newly created rule will depend upon.

To add names to the **Dependent rule names** field, proceed as follows:

1. Locate the dependent rule. You can do this in one of two ways:

### Browse for a rule

If you are familiar with the location of the dependent rule, then click **Browse**. Navigate to the rule's location and highlight it.

### Find a rule

If you are unfamiliar with the location of the rule, then click **Find**. This launches a Find Rules window in which you can specify options and then initiate a search. When you have located the rule, highlight it.

2. Click **Add**.

**5.0.2** To change the name of a rule in the **Dependent rule names** field, highlight one or more rules and click **Edit**. Specify the new name for the rule and click **OK**.

**5.0.2** To change the folder name of a rule in the **Dependent rule names** field, highlight one or more rules and click **Change Folder**. Fill-in the fields in the Change Folder window. All selected rules are updated with the newly specified values.

To delete a rule from the **Dependent rule names** field, highlight the rule and click **Delete**.

To change the order of the rules in the **Dependent rule names** field, highlight a rule and click the up or down arrows to move it to a new location.

## New Rule properties window: Other tab

Use the **New Rule properties window: Other tab** to establish precedence and enter information that is relevant to you, but does not fit into any other category. The following fields are available on this tab:

**Precedence:** Use the **Precedence** field to specify the relative priority when firing the rule. This value is used to order the rules from lowest to highest.

**User defined data:** Use the **User defined data** field to enter any additional text that you want to store. The Business Rule Beans run time does not use this field.

### Import Rules window

Use the Import Rules window to select and import a rule expressed in an XML format. The following fields and options are available in this window:

**File Name:** Use the **File Name** field to specify the name of a file that contains the rules that you want to import. To search for a file, click the **ellipses** icon to the right of the text entry field.

**Show output from rule importer:** Select **Show output from rule importer** to display detailed information about the rules that you want to import.

**Show rules to be created but do not create them:** Select **Show rules to be created but do not create them** to validate the XML prior to committing to the rule's creation. The rule XML runs through the importer but is not created.

**Update existing rules with the same primary key:** Select **Update existing rules with the same primary key** to update the rules with the same primary key. If this check box is clear, the rule is not imported if one is found with this same primary key. An error message is shown and the transaction in which this rule was created is not committed. The primary key is an optional tag within the XML and there is a possibility that a rule already exists on the system with this same primary key.

**Transaction Option:** The following is a list of transaction options:

**One transaction per rule**

Use the **One transaction per rule** option to start a transaction for each rule. If one rule fails to import, it does not prevent other rules in the specified file from being imported.

**One transaction for all rules**

Use the **One transaction for all rules** option to stop all of the rules from being imported if any rule fails to import successfully. Use this feature to ensure that all of your rules are in a consistent state. Typically, it is undesirable to have only a portion of the rules imported successfully since rules might have dependencies on other rules.

### **Select Rules To Export window**

This is the first of three windows in the Export Rules Wizard. Use the Select Rules To Export window to select the rules to export. After entering the appropriate information in the following fields, click **Next**.

**Note:** If you specify a folder, the tool exports the entire contents of the folder including its subfolders.

The following options are available in this window:

**Add:** Use the **Add** option to open a window in which you can type the fully qualified name of a rule or a rule folder. If there are multiple rules with the specified name, they are all exported.

**Find:** Use the **Find** option to open a Find Rules window in which you can specify your search criteria. If there are multiple rules with the same name, only the selected rules are exported.

**Browse button:** Use the **Browse** button to open a window and browse for rules to add to the list. If there are multiple rules with the same name, only the selected rules are exported.

**Remove button:** Use the **Remove** button to remove the selected rules from the list.

**Show output from rule exporter:** Select **Show output from rule exporter** to open a window that contains details about the export operation of the select rules. This window is shown when the export operation begins.

### **Change Effective Dates On Exported Rules window**

This is the second of three windows in the Export Rules Wizard. Use the Change Effective Dates on Exported Rules window to alter the rule's start and end times. This procedure is useful when the application and the rules are tested on a development system prior to being deployed on a production system. You can change the dates of the rules and test on the development system using the current date, even if you plan to use the rules on the production system at a future date.

The rules on the development system are not changed; only the exported version of the rules are changed. After determining whether to select the check boxes in the following descriptions, click **Next**.

The following fields and options are available in this window:

**Change start date and time on exported rules:** Select **Change start date and time on exported rules** to alter the start date of the rules to export. You must specify a valid date and time using the format shown. For example, the format might be "m/d/yy h:mm a" resulting in 10/23/01 1:45 PM.

**Change end date and time on exported rules:** Select **Change end date and time on exported rules** to alter the end date of the rules to export. You must specify a valid date and time using the format shown. For example, the format might be "m/d/yy h:mm a" resulting in 10/23/01 1:45 PM.

### Select File For Rule Export window

This is the third of three windows in the Export Rules Wizard. Use the Select File For Rule Export window to choose a name and location for the exported rule.

1. Browse to an existing directory or create a new one.
2. Type the name of the file. Typically, the file will end with an ".xml" extension.
3. Click **Export**.

### Rule properties window: General tab

Use the **Rule properties window: General** tab to enter general information about the rule. The following fields and options are available in this window:

#### Name and location

##### Folder name

(Required) Use the **Folder name** field to identify the folder in which to create the rule. To browse the existing folders, click the **ellipses** button to the right of the text field.

**Name** (Required) Use the **Name** field to give the rule a name. The name cannot include the forward slash '/' and must contain at least one non-blank character. The name cannot exceed the maximum length of the rule name column in the database table.

#### Period when rule is in effect

##### Start date

(Required) Use the **Start date** field to specify the date and time that the rule goes into effect. If you do not specify the time, a value of midnight is used.

##### End date

Use the **End date** field to determine the date when this rule expires. If you do not specify a value, the rule never expires.

#### Classification

Use the **Classification** section to specify your rule's classification status. Choose one of the following options:

- Rule is not classified and does not perform a classification (default)
- Rule performs a classification
- Rule is classified with the following classification

## Status

### Rule is available for use

Select **Rule is available for use** when the rule is available for use by the Business Rule Beans run-time environment. This feature is useful when you have not finished creating the rule, but you want to save your changes.

### Rule properties window: Implementation tab

Use the **Rule properties window: Implementation** tab to define the manner in which the rule is implemented.

**Note:** The rule contains the data and it does not perform the implementation. Rather, the Java rule implementor implements the rule

For example, suppose you want to create a rule that determines whether a given person is a senior (62 years old or older). To implement this rule, give the "com.ibm.websphere.brb.implementor.RuleGreaterThanEqual" Java rule implementor the value 62 and specify it as an initialization parameter. When the application fires the rule, the person's age is passed to the Java rule implementor as a firing parameter and 62 is passed as an initialization parameter. The person's age is compared against the initialization parameter of 62 and a value of true or false is returned from the Java rule implementor to the application. To change the age at which a person is considered a senior citizen, change the value of the initialization parameter.

The following fields and options are available on this tab:

**Java rule implementor:** (Required) Use the **Java rule implementor** field to specify a class to implement this rule. The initialization and firing parameters that are required are determined by looking at the documentation for the specified Java rule implementor.

**Firing location:** Use the **Firing location** field to determine where the rule is fired. You can fire the Java rule implementor on the server where the rules exist or fire it locally on the client machine. The client can be a servlet running on the server. Specify one of the following values for the firing location:

**Local** Use the **Local** option to fire the Java rule implementor local to the application that fired the rule.

#### **Remote**

Use the **Remote** option to fire the Java rule implementor on the server where the rules exist.

#### **Anywhere (default)**

Use the **Anywhere** option to attempt to fire the Java rule implementor locally first. If the Java rule implementor cannot be found, then it is fired remotely.

To choose the value of the firing location, you must take both performance and maintenance into consideration. Most rules perform better if they are run on the same Java virtual machine (JVM) as the application (locally). However, there might be cases where a Java rule implementor performs server-intensive tasks, in which case the rules might run better when they run on the server. To run locally, you must have all of the Java rule implementors installed locally. They must be accessible by the application that fires the rules.

**Initialization parameters:** The initialization parameters contain constant values passed to the rule implementor when it is initialized. Typically, the initialization parameters contain values that might change as your business practices evolve, such as the age at which a person is considered a senior citizen or the current interest rate for a loan.

To add a new initialization parameter to the list, click **Add** and fill in the fields in the Add Initialization Parameter window.

To edit an existing initialization parameter, highlight it, click **Change**, and fill in the fields in the Change Initialization Parameter window.

To delete an initialization parameter, highlight it and click **Delete**.

To change the order of the initialization parameters, highlight one and click the up or down arrows to move it to a new location.

**Firing parameters:** The firing parameters contain values passed from the trigger point in the application to the Java rule implementor at run time. You can alter the parameters coming from the application before passing them to the Java rule implementor. Typically these parameters are left unaltered.

For example, when implementing the "isSeniorCitizen" rule that determines whether a person is a senior citizen, you might want to pass a person from the application to the rule as the target object. However, the "isSeniorCitizen" rule uses the RuleGreaterThanEqual Java rule implementor, which requires that you pass an integer value. You can alter the firing parameters to specify that the method "getAge" is called on the person object and pass that result to the Java rule implementor.

You must choose one of the following:

**Pass firing parameters from trigger point unchanged**

The parameters specified in the trigger point of the application are passed to the Java rule implementor without being altered. This is the default value.

**Specify firing parameters**

The values specified in the table are passed to the Java rule implementor.

To add a new firing parameter to the list, click **Add** and fill in the fields in the Add Firing Parameter window.

To edit an existing firing parameter, highlight the firing parameter, click **Change**, and fill in the fields in the Change Firing Parameter window.

To delete a parameter, highlight the firing parameter and click **Delete**.

To change the order of the firing parameters, highlight one and click the up or down arrows to move it to a new location.

**Add Firing Parameter window**

Use the Add Firing Parameter window to add a firing parameter to a rule. The firing parameters contain values passed from the trigger point in the application to the Java rule implementor at run time. You can alter the parameters coming from the application before passing them to the Java rule implementor. Typically, these parameters are left unaltered.

To add a firing parameter, enter a field description, select an appropriate option, and click **Add**. The parameter is added and the window remains open to specify additional parameters. When you finish specifying initialization parameters, click **Close**. The following fields and options are available in this window:

**Description:** Use the **Description** field to type a description of the firing parameter.

**Specify a type and value:** Use the **Specify a type and value** option to specify a constant value to pass to the Java rule implementor.

**Get value from method call:** Use the **Get value from method call** option to call the specified method on the target object.

**Get value from trigger point firing parameters:** Use the **Get value from trigger point firing parameters** option to get a specific value from the firing parameters that were specified in the trigger method in the application. An index to the original firing parameter must be specified. This index starts with 0; thus, specify the value 0 to pass the first firing parameter. Specify the trigger point parameter number.

### **Change Firing Parameter window**

Use the Change Firing Parameter window to edit an existing firing parameter. The firing parameters contain values passed from the trigger point in the application to the Java rule implementor at run time. You can alter the parameters coming from the application before passing them to the Java rule implementor. Typically, these parameters are left unaltered.

To change a firing parameter, enter a field description, select an appropriate option, and click **Change**. The parameter is changed and the window remains open to specify additional parameters. When you finish specifying initialization parameters, click **Close**.

The following fields and options are available in this window:

**Description:** Use the **Description** field to type a description of the firing parameter. There are three types of firing parameters that can be specified.

**Specify a type and value:** Use the **Specify a type and value** option to specify a constant value to pass to the Java rule implementor.

**Get value from method call:** Use the **Get value from method call** option to call the specified method on the target object.

**Get value from trigger point firing parameters:** Use the **Get value from trigger point firing parameters** option to pass a specific value from the firing parameters that were specified on the trigger method in the application. An index to the original firing parameter must be specified. This index starts with 0; thus, specify the value 0 to pass the first firing parameter.

### **Rule properties window: Description tab**

Use the **Rule properties window: Description** tab to define the purpose and intent of the rule. All of the fields in this panel are optional and none are used by the Business Rule Beans run-time environment. The following fields are available on this tab:



**Business Intent:** Use the **Business Intent** field to describe the business intent of this rule.

**Description:** Use the **Description** field to define a general description of the rule and its purpose.

**Original requirement:** Use the **Original requirement** field to compose a description of the original requirement that created this rule.

### **Rule properties window: Dependent Rules tab**

Use the **Rule properties window: Dependent Rules** tab to specify the rules that the newly created rule will depend upon.

To add names to the **Dependent rule names** field, proceed as follows:

1. Locate the dependent rule. You can do this in one of two ways:

#### **Browse for a rule**

If you are familiar with the location of the dependent rule, then click **Browse**. Navigate to the rule's location and highlight it.

#### **Find a rule**

If you are unfamiliar with the location of the rule, then click **Find**. This launches a Find Rules window in which you can specify options and then initiate a search. When you have located the rule, highlight it.

2. Click **Add**.

**5.0.2** To change the name of a rule in the **Dependent rule names** field, highlight one or more rules and click **Edit**. Specify the new name for the rule and click **OK**.

**5.0.2** To change the folder name of a rule in the **Dependent rule names** field, highlight one or more rules and click **Change Folder**. Fill-in the fields in the Change Folder window. All selected rules are updated with the newly specified values.

To delete a rule from the **Dependent rule names** field, highlight the rule and click **Delete**.

To change the order of the rules in the **Dependent rule names** field, highlight a rule and click the up or down arrows to move it to a new location.

### **Change Folder window**

Use the Change Folder window to change the folder name for all selected dependent rules. There are two options for changing the folder name:

**Specify a new folder name:** Use the **Name** field to specify the new folder name for all selected rules. All selected dependent rule names are updated to contain the new folder, exactly as specified.

**Replace common portion of the folder name:** Use this option to change the portion of the folder name that is common to all selected rules.

For example, suppose the following two rules exist:

- com/acme/insurance/health/ruleA
- com/acme/insurance/life/ruleB

You want to change the `com/acme/insurance` portion of the name to `com/acme/ins`. The `com/acme/insurance` portion will be changed while leaving the remaining portion of the rule name intact.

The **Common folder** field is set to the common portion of the folder name. In the preceding example, this field is set to `com/acme/insurance`. This value cannot be changed.

Specify the new folder name in the **Replace with** field. For the preceding example, specify `com/acme/ins`.

### **Rule properties window: Other tab**

Use the **Rule properties window: Other** tab to establish precedence and enter information that is relevant to you, but does not fit into any other category. The following fields are available in this window:

**Precedence:** Use the **Precedence** field to specify the relative priority when firing the rule. This value is used to order the rules from lowest to highest.

**User defined data:** Use the **User defined data** field to enter any additional text that you want to store. The Business Rule Beans run-time environment does not use this field.

### **Validate Rules**

Use the **Validate** option to verify the correctness of the selected rules or folder. A report is generated and displayed in system's browser window. You can use the browser functionality to view, save, and print the results. The following items are validated:

#### **Start and End dates**

A valid start date must be specified. The end date must be empty (never expires) or be later than the start date.

#### **Java rule implementor name**

A valid and existing Java rule implementor name must be specified. The specified class must implement the `RuleImplementor` interface.

#### **Dependent rules**

The existence of each specified dependent rule is verified.

You can specify a Java classpath to use to validate rule implementor classes for the rules. To add individual JAR files or directories containing classes to the **Classpath** field, click **Add** or **Browse**.

To specify the full name of a jar file or directory to be added to the classpath, click **Add**.

To search through the file system for a JAR file or directory to be added to the classpath, click **Browse**.

To change an existing classpath entry, highlight the entry and click **Edit**. Specify the new name for the classpath entry and click **OK**.

To remove a classpath entry from the list, highlight one or more entries and click **Remove**.

To change the order of the classpath entries, highlight one or more entries and click the Up or Down arrows to move it to a new location.

After specifying the classpath, click **OK** to generate the report.

## Edit menu

This menu describes the options available on the Edit menu window. The following options are available in this window:

### Cut

Use the **Cut** option to move rules and rule folders.

### Copy

Use the **Copy** option to copy rules and rule folders. The following tasks can be accomplished using the **Copy** option:

- **Copying rules**

Copy or move a rule from one folder to another by either cutting and pasting it or dragging and dropping it.

#### Cutting and pasting

Use menu commands (**Edit > Copy**, **Edit > Cut** and **Edit > Paste**) or keyboard commands (**CTRL-C**, **CTRL-V** and **CTRL-X**) to copy or move a rule.

#### Dragging and dropping

Highlight the rule you want to copy. Then, press and hold the right mouse button, drag the cursor to the target location, and release. Select **Copy** or **Move** from the list.

**Note:** You also can copy a rule so that the copy replaces the existing rule at a specified date. This is referred to as a Quick Copy.

- **Copying rule folders**

Copy or move a rule folder and all its contents by either cutting and pasting it or dragging and dropping it.

#### Cutting and pasting

Use menu commands (**Edit > Copy**, **Edit > Cut** and **Edit > Paste**) or keyboard commands (**CTRL-C**, **CTRL-V** and **CTRL-X**) to copy or move a rule.

#### Dragging and dropping

Press and hold the right mouse button on the folder to be copied. Drag the cursor to the target location and release the mouse button. Select **Copy** or **Move** from the list.

### Paste

Use the **Paste** option to add cut or copied rules and rule folders.

### Find

Use the **File** option to search for a rule. A window opens in which you can specify your search criteria. If you would like to search a specific folder, right-click the folder and select **Find** from the list. For more information, click the **Find Rules window** link under Related reference.

- **Search the whole directory**

To search the whole directory, complete the following steps:

1. Click **Edit > Find** from the main menu of the Rule Browser.
2. Determine your search criteria in the Find Rules window.

- **Search a specific folder**

To search a specific folder, complete the following steps:

1. Right-click the folder and select **Find** from the list.
2. Determine your search criteria in the Find Rules window.

The results of your search are displayed in a Search Results window.

### **Quick Copy**

Use the **Quick Copy** option to make a copy of a rule that will replace the existing one on a specified date. You can modify the copy so that a new value goes into effect on the desired date and time while the old rule expires. For more information, see the **Quick Copy window** link in Related reference.

### **Select All**

Use the **Select All** option to facilitate rule selection.

### **Deselect All**

Use the **Deselect All** option to deselect rules.

### **Quick Copy window**

Use the Quick Copy window to make a copy of a rule that replaces the existing one on a specified date.

For example, suppose that you have an "isSeniorCitizen" rule. Currently, a person is considered a senior citizen if they are 62 years of age or older. Starting on January 1, 2002, you must change this to 65. Use Quick Copy to specify the new date and to change the age from 62 to 65. The current rule is set to expire on the same date that the new rule takes effect. The new senior citizen age is defined as 65.

**Note:** Use the **Quick Copy** function for simple changes only.

In the following fields, specify how the copy differs from the original:

**Start Date For New Rule:** Use the **Start Date For New Rule** field to enter the date that the new rule replaces the existing rule. Use the date and time format that is shown. For example, the format might be the following:

Usage scenario

M/d/yy h:mm a.

**Change parameter values for new rule:** Use the **Change parameter values for new rule** field to add new parameter values.

### **Automatically Update References window**

Use the Automatically Update References window to indicate what you would like to do with rules that contain references to the rules being moved, copied, or renamed. You can choose whether to update references and whether to view the references before updating the rules. When the rule references are displayed, you can choose which references to update. The following options are available in this window:

**Automatically update references:** Select **Automatically update references** to indicate that references should be updated.

**Display the updates before changing the rules option:** Select **Display the updates before changing the rules option** to indicate whether to show the changes

before updating the references. When you choose this option, the Update Dependent Rule References window is displayed, allowing you to select which references to update.

### **Update Dependent Rule References window**

Use the Update Dependent Rule References window to select the dependent rule references that must be updated as a result of the move, copy, or rename operation.

Each row in the table represents one dependent rule reference. The following information is shown:

**Rule Name**

Specifies the name of the rule containing the dependent rule reference.

**Folder Name**

Specifies the folder containing the rule specified by **Rule Name**.

**Dependent Rule Old Name**

Specifies the name of the dependent rule that is to be changed as a result of this move, copy, or rename operation.

**Dependent Rule New Name**

Specifies the new name to be given to the dependent rule as a result of this move, copy, or rename operation.

**Status** Specifies the status of the rule; in effect, scheduled, expired, or invalid.

**Start Date**

Specifies the date and time when this rule will begin being used.

**End Date**

Specifies the date and time of when this rule will expire (no longer be used).

The check box to the left of the rule name indicates whether the dependent rule reference will be updated. Select the check box to change the selection which indicates that the reference should be updated.

Use the **Select All** button to indicate that all dependent rule references should be updated.

Use the **Deselect All** button to clear the dependent rule references selection.

Click **Continue** after reviewing the list of dependent rule references and indicating which references to update. Click **Cancel** to cause the move, copy, or rename operation to be cancelled completely.

### **References To Rules Being Deleted window**

Use the References To Rules Being Deleted window to view dependent rule references when a rule or folder is deleted. Each row in the table represents one dependent rule reference. The following information is shown:

**Rule Name**

Specifies the name of the rule containing the dependent rule reference.

**Folder Name**

Specifies the name of the folder containing business rules.

**Dependent Rule Reference**

Specifies the name of the dependent rule that is being deleted.

**Start Date**

Specifies the date and time when this rule is implemented.

**End Date**

Specifies the date and time when this rule expires.

Click **Continue** to delete. Click **Cancel** to abort.

**Rule Names Already Exist window**

The Rule Names Already Exist window displays when one or more rules are copied into a folder and existing rules are found with the same name. This window enables you to rename the rules before you proceed with the copy operation. The conflicting rule names do not need to be changed; you can create multiple rules with the same name.

Each line of the table represents one rule for which a name conflict has occurred. The table shows the following fields:

**Rule Name**

Specifies the name of the rule being copied which is the same as the name of a rule that already exists in the destination folder.

**Folder Name**

Specifies the folder containing the rule specified by Rule Name

**Status** Specifies the status of the rule; in effect, scheduled, expired, or invalid

**Start Date**

Specifies the date and time when this rule will begin being used

**End Date**

Specifies the date and time when this rule will expire (no longer be used)

To change the name of the rule being copied, click on the Rule Name in the table and specify a new name for the rule.

Click **OK** when you are finished reviewing or changing the rule names. Click **Cancel** to abort the copy operation.

## View menu

This article describes the options available in the View menu window. The following options are available in this window:

**Status Bar**

Use the **Status Bar** option to toggle the status bar on or off. The status bar is shown along the bottom of the Rule Browser window.

**Specify Columns**

Use the **Specify Columns** option to adjust the type and order of the columns that display in your window.

In the window that opens, the following tasks can be accomplished:

**Add a new column**

Select one or more entries in the **Available columns** list and click the **Add** button. The selected entries are added to the end of the **Columns displayed** list.

### Remove a column

Select one or more entries in the **Columns displayed** and click the **Remove** button. The selected entries are added to the end of the **Available columns** list.

### Reorder columns

Select one or more entries in the **Columns displayed** list. To move the entries towards the top of the list, click the Up arrow; to move the entries towards the bottom, click the Down arrow.

### Specify Date/Time Format

Use the **Specify Date/Time Format** option to adjust the format used when displaying dates and times. For more information, click the **Specify Date/Time Format window** link under Related reference.

### Refresh

Use the **Refresh** option to update the contents of the folder hierarchy and the rule table.

### Specify Date/Time Format window

Use the Specify Date/Time Format window to change the date and time format. Choose one of the following radio button options and then click **OK**:

**Use default format for this locale:** Use the **Use default format for this locale** option to adjust the date and time format to match the default setting of your current locale.

**Select a predefined format for the date and time:** Use the **Select a predefined format for the date and time** option to select one of several existing date and time formats.

**Specify a custom format for the date and time:** Use the **Specify a custom format for the date and time** option to determine your own format for your date and time display. Choose one of the date and time formats from the two menus.

## Find Rules window

Use the Find Rules window to specify search criteria to locate rules. The search combines your queries using a logical "AND" operation. For example, if you specify both a folder name and a rule name, the search finds rules that match both the folder name and rule name. It displays the results in a Search Results window.

**Note:** All of the fields in this window are optional.

### Main menu

The main menu has the following options:

- **File**

- **Save As**

- Use the **Save as** option to open a Save Search window and store the current search criteria.

- **Open**

- Use the **Open** option to open the Open Saved Search window and load a previously saved set of search criteria into the Find Rules window.

- **Close**

- Use the **Close** option to close the Find Rules window.

- **View**

### Show Search

Use the **Show Search** option to display a text description of your search criteria on one screen.

### Specify Date/Time Format

Use the **Specify Date/Time Format** option to adjust the format used when displaying dates and times. For more information, see the **Specify Date/Time Format window** link in Related reference.

## Tabs

The following tabs are displayed in this window:

**Name** Use the **Name** tab to specify the basic search criteria. For more information, click the **Find Rules window: Name tab** link under Related reference.

**Date** Use the **Date** tab to specify the date-related search criteria. For more information, click the **Find Rules window: Date tab** link under Related reference.

### Classification

Use the **Classification** tab to specify the search criteria related to a rule's classification. For more information, click the **Find Rules window: Classification tab** link under Related reference.

### Implementation

Use the **Implementation** tab to specify the search criteria that is based on the manner in which the rule is implemented. For more information, click the **Find Rules window: Implementation tab** link under Related reference.

### Description

Use the **Description** tab to specify the text-based search criteria related to a rule's description. For more information, click the **Find Rules window: Description tab** link under Related reference.

**Other** Use the **Other** tab to specify the search criteria based on precedence and user-defined data. For more information, click the **Find Rules window: Other tab** link under Related reference.

## Find Rules window: Name tab

Use the **Find Rules window: Name** tab to specify the basic search criteria. The following fields and options are available on this tab:

**Name:** Use the **Name** field to search for a specific rule name.

**Note:** This is case-sensitive.

**Drop-down search option list:** Use the one of the following options in the **Drop-down search option list** to narrow your search:

**equal** Use the **equal** selection to look for an *exact* match.

### starting with

Use the **starting with** selection to find rules whose name *starts* with the specified value.

### ending with

Use the **ending with** selection to find rules whose name *ends* with the specified value.



**containing**

Use the **containing** selection to find rules that *contain* the specified value.

**Location:** Use the **Location** field to specify the folder that you want to search. Click the **ellipses** button to the right of the field if you want to browse for the folder.

**Note:** The folder names are case-sensitive.

**Include subfolders:** Select **Include subfolders** to include the folder's subfolders in the search.

**Status:** Use the **Status** menu to specify a search criteria that is based on a rule's availability.

**Find Rules window: Date tab**

Use the **Find Rules window: Date** tab to specify date-related search criteria. A rule always has a start date and a range of time in which it is in effect. The end date is optional and if it is not specified, the rule never expires. The following options are available on this tab:

**Find Rules for any date:** Use the **Find Rules for any date** option to remove the date from consideration in the search criteria.

**Find Rules that are:** Use the **Find Rules that are** option to search for rules in one of the following states:

**currently in effect**

The **currently in effect** selection finds rules that are active at this point in time.

**scheduled**

The **scheduled** selection finds rules that go into effect at a future date.

**expired**

The **expired** selection finds previously active rules that are beyond the rule's end date.

**Find Rules:** Use the **Find Rules** option to specify the dates you want to search. Modify the following criteria to narrow your search:

1. Select one of the following rule states from the menu:

- in effect
- starting
- ending

2. Select **Query on date only** if you would like your search to ignore time-specific information.

3. Select one of the following three methods by which to search:

**on x** Select the **on x** option to find rules that are in the desired state (as chosen in Step 1) on the specified date. If the **Query on date only** check box is clear, then midnight is used for the time.

**anytime between x and y**

Select the **anytime between x and y** option to find rules that are in the desired state (as chosen in Step 1) *anytime* between the given dates. Specify a start date (represented by *x*) and an end date (represented by *y*). If the **Query on date only** check box is clear, then midnight of each day is used for the time.

### **anytime during the next $x$ days**

Select the **anytime during the next  $x$  days** option to find rules that are in the desired state (as chosen in Step 1) within this period of time and specified in days.

### **Find Rules window: Classification tab**

Use the **Find Rules window: Classification** tab to specify search criteria related to a rule's classification. The following options are available on this tab:

**Show all:** Use the **Show all** option if you do not want to include the classification information in the search criteria.

**Show rules that are not classified and do not perform classification:** Use the **Show rules that are not classified and do not perform classification** option to find rules that you do not need to classify.

**Show rules that perform classification:** Use the **Show rules that perform classification** option to find rules that return a classification such as "Gold", "Silver", or "Bronze".

**Show rules that are classified:** Use the **Show rules that are classified** option to find rules that are classified with the specified classification. The specified classification is case-sensitive. Enter a specific classification into the field that is provided.

### **Find Rules window: Implementation tab**

Use the **Find Rules window: Implementation** tab to specify search criteria that is based on the manner in which the rule is implemented. The following options are available on this tab:

**Java rule implementor:** Use the **Java rule implementor** option to search for rules that use the specified Java rule implementor. You can use one of the values in the list or type in your own.

**Firing location:** Use the **Firing location** option to search for the location in which the rule implementor is run. Choose one of the following values from the check box:

**Local** Use the **Local** option to search locally for the location in which the rule implementor is run.

#### **Remote**

Use the **Remote** option to search the server for the location in which the rule implementor is run.

#### **Anywhere**

Use the **Anywhere** option to search both locally and on the server for the location in which the rule implementor is run.

**Firing parameters:** Use the **Firing parameters** option to search for rules that alter the firing parameters passed from the trigger point to the Java rule implementor. For more information, see the **Add Firing Parameter window** link in Related reference.

The following selections are available for the firing parameter option:

- show rules that alter firing parameters
- show rules that do not alter firing parameters

### Find Rules window: Description tab

Use the **Find Rules window: Description** tab to specify text-based search criteria related to a rule's description. The following fields and options are available on this tab:

**Business intent:** Use the **Business intent** field to search for rules with a given business intent.

**Description:** Use the **Description** field to search for rules with a given description.

**Original requirement:** Use the **Original requirement** field to search for rules with a given original requirement.

**Drop-down search option list:** Use the following options in the **Drop-down search option list** to narrow your search in the **Business intent**, **Description**, and **Original requirements** fields:

**equal** Use the **equal** selection to look for an *exact* match.

**starting with**

Use the **starting with** selection to find rules whose name *starts* with the specified value.

**ending with**

Use the **ending with** selection to find rules whose name *ends* with the specified value.

**containing**

Use the **containing** selection to find rules that *contain* the specified value.

### Find Rules window: Other tab

Use the **Find Rules window: Other** tab to specify search criteria based on precedence and user-defined data. The following options are available on this tab:

**Precedence:** Use the **Precedence** option to search for rules with given precedence. The precedence is an integer value that specifies the relative priority of this rule when it is fired.

**Drop-down search option list:** Use the **Drop-down search option list** to narrow your search by using one of the following options:

- equal
- less than
- less than or equal
- greater than
- greater than or equal
- not equal

**User defined data:** Use the **User defined data** option to search for rules with given user defined data.

**Drop-down search option list:** Use one of the following options from the **Drop-down search option list**:

**equal** Use the **equal** selection to look for an *exact* match.

**starting with**

Use the **starting with** selection to find rules whose name *starts* with the specified value.

**ending with**

Use the **ending with** selection to find rules whose name *ends* with the specified value.

**containing**

Use the **containing** selection to find rules that *contain* the specified value.

**Search Results window**

This window contains the results of a search from a Find Rules window. It is nearly identical to the Rule Browser in terms of form and function. Use the options in the main menu to perform many of the same administrative actions, with a few differences:

- Rules cannot be imported into the Search Results window.
- Rules cannot be pasted into the Search Results window.
- The **Open Containing Folder** menu item opens a new browser window showing the contents of the folder for the selected rule.

**Save Search window**

Use the Save Search window to store the current search criteria for later retrieval. Follow the following steps to save your search:

1. Type in a name for your search or replace a previously saved search from the list.
2. Click **OK**.

The saved search criteria is loaded into the Find Rules window.

**Open Saved Search window**

Use the Open Saved Search window to open a previously saved search. Follow these steps to open your saved search:

1. In the Select a Search window, highlight the name of the search you want to open.
2. Click **OK**.

The saved search criteria is loaded into the Find Rules window.

---

## Business rule beans: Resources for learning

Use the following links to find relevant supplemental information about business rule beans. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Planning, business scenarios, and IT architecture
- Programming instructions and examples
- Administration

## Planning, business scenarios, and IT architecture

- ✓ **Delivering new business value to the enterprise on a J2EE and Web services base (Update)**  
(<http://www7b.software.ibm.com/wsdd/library/summaries/200462.html>)  
This paper, in PDF form, describes the strategy behind the IBM extensions to J2EE and Web services functionality in the WebSphere Application Server Version 4.0 Enterprise Edition. It explains Enterprise Services, business rule beans, message beans and JMS listener, internationalization, shared work areas, bidirectional CORBA connectivity, the ActiveX bridge, and the C++ CORBA SDK.
- 📄 **WebSphere Application Server V5.0 Architecture and Overview**  
([http://developerworks.cybercentral.com/ibm0502/amt/ibmpresentations/683\\_1.pdf](http://developerworks.cybercentral.com/ibm0502/amt/ibmpresentations/683_1.pdf))  
This is an IBM developerWorks presentation that provides an overview of the functionality available with WebSphere Application Server V5.0.

## Programming instructions and examples

- 📄 **Message-Driven Beans and Encapsulated Business Rules**  
(<http://www2.theserverside.com/resources/article.jsp?l=Message-Driven-Beans-And-Encapsulated-Business-Rules>)  
This article describes how to use business rules with Message-driven Beans.
- 📄 **WebSphere Application Server Enterprise Edition 4.0: A Programmer's Guide** (<http://www.redbooks.ibm.com/redbooks/SG246504.html>)  
Chapter three of this programmer's guide provides information about implementation, modification, and deployment of business rules.
- 📄 **WebSphere Application Server Enterprise Edition Technology Sample**  
([http://www7b.software.ibm.com/wsdd/downloads/ee41\\_landing.html](http://www7b.software.ibm.com/wsdd/downloads/ee41_landing.html))  
This technology sample enables developers to gain experience with the Business Rule Beans technology.

## Administration

- 📄 **IBM WebSphere Administration** (<http://books.mcgraw-hill.com/cgi-bin/pbg/0072223154.html>)

•



### Listing of all IBM WebSphere Application Server Redbooks

(<http://publib-b.boulder.ibm.com/Redbooks.nsf/Portals/WebSphere>)

This is a listing of the Redbook publications about the WebSphere software platform.

- 📄 **WebSphere Application Server Version 4.0 Enterprise Edition — Presentations and Labs** ([http://www7b.boulder.ibm.com/wsdd/library/presents/WAS\\_EE\\_Training.html](http://www7b.boulder.ibm.com/wsdd/library/presents/WAS_EE_Training.html))



---

## Chapter 17. Using asynchronous beans

The asynchronous beans feature adds a new set of APIs that enable J2EE applications to run work asynchronously inside a WebSphere Enterprise Application Server. This topic provides a brief overview of the tasks involved in using asynchronous beans. For a more detailed description of the asynchronous beans model, review the conceptual topic [Asynchronous beans](#).

### Steps for this task

1. Configuring work managers
2. Assembling applications that use work managers
3. Developing work objects to run code in parallel
4. Developing event listeners
5. Developing asynchronous scopes

---

## Asynchronous beans

An asynchronous bean is a Java object or enterprise bean that can be executed asynchronously by a J2EE application, using the J2EE context of the bean's creator.

Asynchronous beans can improve performance by enabling a J2EE program to decompose operations into parallel tasks. Asynchronous beans enable the construction of stateful, "active" J2EE applications. These applications address a segment of the application space that J2EE has not previously addressed (that is, advanced applications that require application threading, active agents within a server application, or distributed monitoring capabilities).

Asynchronous beans can run using the J2EE security context of the creator J2EE component. These beans also can run with copies of other J2EE contexts. For example:

- Internationalization context
- Application profiles
- Work areas
- Access intent policies

### **Asynchronous bean interfaces**

There are three types of asynchronous bean:

#### **Work object**

A work object implements the `com.ibm.websphere.asynchbeans.Work` interface. A work object runs parallel to its caller using the `WorkManager.startWork()` method. Applications implement work objects in order to run code blocks asynchronously. For more information on the `Work` interface, see the [Related reference](#) section at the end of this article.

#### **Alarm listener**

An alarm listener is an object that implements the `com.ibm.websphere.asynchbeans.AlarmListener` interface. Alarm listeners

are called when a high-speed transient alarm expires. For more information on the AlarmListener interface, see the Related reference section at the end of this article.

### **Event listener**

An event listener can implement any interface. An event listener is a lightweight, asynchronous notification mechanism for asynchronous events within a single JVM. An event listener would typically be used to enable J2EE components within a single application to notify each other about various asynchronous events.

### **Supporting interfaces**

#### **Work manager**

A work manager is a thread pool that administrators create for J2EE applications. The administrator specifies the properties of the thread pool and a policy that determines which J2EE contexts the asynchronous bean will inherit.

#### **Event source**

An event source implements the `com.ibm.websphere.asynchbeans.EventSource` interface. An event source is a system-provided object that supports a generic, type-safe asynchronous notification server within a single JVM. The event source enables event listener objects, which implement any interface, to be registered. For more information on the EventSource interface, see the Related reference section at the end of this article.

#### **Event source events**

Every event source can generate events of its own. Event sources also can generate their own events such as 'listener count changed'. An application can register an event listener object that implements `com.ibm.websphere.asynchbeans.EventSourceEvents`. This enables the application to catch events such as listeners being added or removed, or a listener throwing an unexpected exception. For more information on EventSourceEvents, see the Related reference section at the end of this article.

Additional interfaces, including alarms and subsystem monitors, are introduced in the topic *Developing Asynchronous scopes*, which discusses some of the advanced applications of asynchronous beans.

### **Transactions**

Every asynchronous bean method is called using its own transaction, much like container-managed transactions in a typical enterprise bean. It is very similar to the situation when an EJB method is called with `TX_NOT_SUPPORTED`. The run-time environment starts a local transaction before invoking the method. The asynchronous bean method is free to start its own global transaction if this is possible for the calling J2EE component. For example, if an enterprise bean creates the component, the method that creates the asynchronous bean must be `TX_BEAN_MANAGED`.

When you call an entity bean from within an asynchronous bean, for example, you must have a global transactional context available on the current thread. Since asynchronous bean objects start local transactional contexts, you can encapsulate all entity bean logic in a session bean that has a method marked as `TX_REQUIRES` or



equivalent. This process establishes a global transactional context from which you can access one or more entity bean methods.

If the asynchronous bean method throws an exception, any local transactions are rolled back. If the method returns normally, any incomplete local transactions are completed according to the unresolved action policy configured for the bean. EJB methods can configure this policy using their deployment descriptor. If the asynchronous bean method starts its own global transaction and does not commit this global transaction, the transaction is rolled back when the method returns.

### Access to J2EE component meta-data

If an asynchronous bean is a J2EE component, such as a session bean, its own meta-data is active when a method is called. If an asynchronous bean is a simple Java object, the J2EE component metadata of the creating component is available to the bean. Like its creator, the asynchronous bean can look up the `java:comp` namespace. This enables the bean to access connection factories and enterprise beans, just as it would if it were any other J2EE component. The environment properties of the creating component also are available to the asynchronous bean.

The `java:comp` namespace is identical to the one available for the creating component; the same restrictions apply. For example, if the Enterprise Bean or servlet has an EJB reference of `java:comp/env/ejb/MyEJB`, this EJB reference is available to the Asynchronous Bean. In addition, all of the connection factories use the same resource-sharing scope as the creating component.

### Connection management

An asynchronous bean method can use the connections that its creating J2EE component obtained using `java:comp` resource references. (For more information on resource references, see [References](#)). However, the bean method must access those connections using a `get`, `use`, `close` pattern. There is no connection caching between method calls on an asynchronous bean. The connection factories or `DataSources` themselves can be cached, but the connections must be retrieved on every method call, used, and then closed. While the asynchronous bean method can look up connection factories using a global JNDI name, this is not recommended for the following reasons:

- The JNDI name is hard-coded in the application (for example, as a property or string literal).
- The connection factories are unshared because there is no way to specify a sharing scope.

For code examples that demonstrate both the correct and the incorrect ways to access connections from asynchronous bean methods, see the topic [Example: Asynchronous bean connection management](#).

## Example: Asynchronous bean connection management

An asynchronous bean method can use the connections that its creating J2EE component obtained using `java:comp` resource references. (For more information on resource references, see the topic [References](#).) The following is an example of an asynchronous bean that uses connections correctly:

```
class GoodAsynchBean
{
    DataSource ds;
    public GoodAsynchBean()
```

```

throws NamingException
{
// ok to cache a connection factory or datasource
// as class instance data.
InitialContext ic = new InitialContext();
// we are assuming that the creating J2EE component has this
// resource reference defined in its deployment descriptor.
ds = (DataSource)ic.lookup("java:comp/env/jdbc/myDataSource");
}
// When the asynchronous bean method is called, get a connection,
// use it, then close it.
void anEventListener()
{
Connection c = null;
try
{
c = ds.getConnection();
// use the connection now...
}
finally
{
if(c != null) c.close();
}
}
}

```

The following is an example of an asynchronous bean that uses connections illegally:

```

class BadAsynchBean
{
DataSource ds;
// Do not do this. You cannot cache connections across asynch method calls.
Connection c;

public BadAsynchBean()
throws NamingException
{
// ok to cache a connection factory or datasource as
// class instance data.
InitialContext ic = new InitialContext();
ds = (DataSource)ic.lookup("java:comp/env/jdbc/myDataSource");
// here, you broke the rules...
c = ds.getConnection();
}
// Now when the asynch method is called, illegally use the cached connection
// and you'll likely see a bunch of J2C related exceptions at runtime.
// close it.
void someAsynchMethod()
{
// use the connection now...
}
}

```

---

## Configuring work managers

### Before you begin

If you are unfamiliar with work managers, review the conceptual topic [Work managers](#).

A work manager acts as a thread pool for application components that use asynchronous beans. Use the administrative console to configure work managers. You can define multiple work managers for each cell. Each work manager is bound to a unique place in JNDI.

**Note:** The work manager service is only supported from within the EJB Container or Web Container (EJBs or Servlets). Looking-up and using a configured WorkManager from a J2EE Application Client container is not supported.

#### Steps for this task

1. Start the administrative console.
2. Select **Resources > Work Managers**.
3. Click **New**.
4. Specify the following required properties:

**Name** The display name for the work manager.

#### **JNDI Name**

The JNDI name for the work manager. This name is used by asynchronous beans that need to look up the work manager. Each work manager must have a unique JNDI name within the cell.

#### **Number of Alarm Threads**

The maximum number of threads to be used for processing alarms. There is a single thread that is used to monitor pending alarms and dispatch them. Also there is an additional pool of threads that is used for dispatching the threads. All alarm managers on the asynchronous beans associated with this work manager share this set of threads. That is, there is a single alarm thread pool for each work manager and all of the asynchronous beans associated with the work manager share this pool of threads.

#### **Minimum Number Of Threads**

The initial number of threads to be created in the thread pool.

#### **Maximum Number Of Threads**

The maximum number of threads to be created in the thread pool. The maximum number of threads can be exceeded temporarily if the **Growable** checkbox is selected. These additional threads are discarded when the work on the thread completes.

#### **Thread Priority**

The order of priority for threads available in the thread pool.

5. [Optional] Specify a **Description** and a **Category** for the work manager.
6. [Optional] Select the **Service Names** (J2EE contexts) on which you want this work manager to be made available. Any asynchronous beans that use this work manager then will inherit the selected J2EE contexts from the component that creates the bean. The list of selected services also is known as the "sticky" context policy for the work manager.  
**Note:** Selecting more services than are actually required might impede performance.
7. Save your configuration.

#### Results

The work manager is now configured and ready to be accessed by application components that need to manage asynchronous code execution.

## **Work managers**

A work manager is a thread pool created for J2EE applications that use asynchronous beans.

Using the administrative console, an administrator can configure any number of work managers. The administrator specifies the properties of the work manager, including the "sticky" context (inheritance) policy for any asynchronous beans that use the work manager. The administrator binds each work manager to a unique place in JNDI.

When writing a Web or EJB component that uses asynchronous beans, the developer should include a resource reference in each component that needs access to a work manager. (For more information on resource references, see the topic *References*.) The component looks up a work manager using a logical name in the component's `java:comp` namespace, just as it would look up a `datasource`, `enterprise bean`, or `connection factory`.

The deployer binds physical work managers to logical work managers when the application is deployed.

For example, if a developer needs three thread pools to partition work between bronze, silver, and gold levels, the developer writes the component to pick a logical pool based on an attribute in the client application's profile. The deployer has the flexibility to decide how to map this request for three thread pools. The deployer might decide to use a single thread pool on a small machine. In this case, the deployer binds all three resource references to the same work manager instance (that is, the same JNDI name). A larger machine might allow for three thread pools, so the deployer binds each resource reference to a different work manager. Work managers can be shared between multiple J2EE applications installed on the same server.

An application developer can use as many logical work managers as necessary; the deployer chooses whether to map one physical work manager or several to the logical work manager defined in the application.

**Note:** All J2EE components that need to share asynchronous scope objects must use the same work manager. These scope objects have an affinity with a single work manager so an application that uses `AsynchScopes` should verify that all of the components using scope objects use the same work manager.

When multiple work managers are defined, the underlying thread pools are created in a JVM only if an application within that JVM looks up the work manager. For example, there might be ten thread pools (work managers) defined, but none are actually created until an application looks them up.

### **How to look up a work manager**

An application can look up a work manager as follows. Here, the component contains a resource reference named `wm/myWorkManager`, which was bound to a physical work manager when the component was deployed:

```
InitialContext ic = new InitialContext();
WorkManager wm = (WorkManager)ic.lookup("java:comp/env/wm/myWorkManager");
```

### **"Sticky" J2EE contexts**

Asynchronous beans can inherit the following J2EE contexts. In other words, these contexts can be made "sticky":

#### **Internationalization context**

#### **Work area**

## Application profile

### Security

The asynchronous bean can be run as anonymous or as the client authenticated on the thread that created it. This is useful because the asynchronous bean can do only what the caller can do. This is more useful than a RUN\_AS mechanism, for example, which prevents this kind of behavior.

### Component meta-data

Component meta-data is relevant only when the asynchronous bean is a simple Java object. If the bean is a J2EE component, such as an enterprise bean, the component's meta-data is active.

Which contexts are sticky depends on the work manager used by the application that creates the asynchronous bean. Using the administrative console, the administrator defines the sticky context policy of a work manager by selecting the services on which the work manager is to be made available.

## Work manager collection

Use this page to view the configuration properties of work managers.

A work manager contains a pool of threads bound into JNDI.

To view this administrative console page, click **Resources > Work Managers** .

### Name

The name by which the work manager is known for administrative purposes.

Data type String

### JNDI Name

The JNDI name used to look up the work manager in the namespace.

Data type String

### Description

A description of this work manager for administrative purposes.

Data type String

### Category

A string that can be used to classify or group this work manager.

Data type String

### Number of Alarm Threads

The number of threads used to execute concurrent alarms.

Data type Integer

### Minimum Number of Threads

The minimum number of threads available in this work manager for running works.

Data type Integer

### Maximum Number of Threads

The maximum number of threads available in this work manager for running works.

Data type Integer

### Thread Priority

The priority of the threads available in this work manager

Data type Integer

### Growable

Specifies whether the number of threads in this work manager can be increased.

### Service Names

A list of service names on which this work manager is made available.

The context information of each selected service is propagated to each work or alarm that is created using this work manager. Selecting services that are not needed can negatively impact performance.

### Work manager settings

Use this page to modify work manager settings.

A work manager contains a pool of threads bound into JNDI.

To view this administrative console page, click **Resources > Work Managers > *workmanager\_name***.

**Name:** The name by which the work manager is known for administrative purposes.

Data type String

**JNDI Name:** The JNDI name used to look up the work manager in the namespace.

Data type String

**Description:** A description of this work manager for administrative purposes.

Data type String

**Category:** A string that can be used to classify or group this work manager.

Data type String

**Number of Alarm Threads:** The number of threads used to execute concurrent alarms.

Data type Integer

**Minimum Number of Threads:** The minimum number of threads available in this work manager for running works.

Data type Integer

**Maximum Number of Threads:** The maximum number of threads available in this work manager for running works.

Data type Integer

**Thread Priority:** The priority of the threads available in this work manager

Data type Integer

**Growable:** Specifies whether the number of threads in this work manager can be increased.

**Service Names:** A list of service names on which this work manager is made available.

The context information of each selected service is propagated to each work or alarm that is created using this work manager. Selecting services that are not needed can negatively impact performance.

## Work manager service settings

Use this page to enable or disable the work manager service, which manages work manager resources used by the server.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Work Manager Service** .

### Startup

Specifies whether the server will attempt to start the work manager service.

Default Selected

Range

#### Selected

When the application server starts, it attempts to start the work manager service automatically.

#### Cleared

The server does not try to start the work manager service. If work manager resources are to be used on this server, the system administrator must start the work manager service manually or select this property then restart the server.

---

## Assembling applications that use work managers

Before you begin

Your administrator needs to configure at least one work manager using the administrative console.

If your application references one or more logical work managers, the logical work managers must be bound to one or more physical work managers using the Application Assembly Tool.

#### Steps for this task

1. Declare a resource reference for each work manager (required action by the application developer). This forms an EAR file. (For more information on resource references, see the topic References.)
2. Using the Application Assembly Tool (AAT), bind each resource reference to a physical work manager.
3. Add a resource reference with the type 'com.ibm.websphere.asynchbeans.WorkManager' to the application's descriptor. The application then can look up this work manager using its resource reference name in java:comp. The AAT or WebSphere Studio Application Developer Integration Edition (WSAD-IE) then can specify which resource references are bound to a physical work manager.

**Note:** The previous process is the same as the process used for DataSources.

---

## Developing work objects to run code in parallel

#### Before you begin

Your administrator must have configured at least one work manager using the administrative console.

To run code in parallel, or in a different J2EE context, wrap the code in a work object.

#### Steps for this task

1. Create a work object.  
A work object implements the com.ibm.websphere.asynchbeans.Work interface.  
For example:  

```
class SampleWork implements Work
```
2. Determine the number of work managers needed by this application component.
3. Look up the work manager or managers using the work manager's resource reference (or logical name) in the java:comp namespace. (For more information on resource references, see the topic References.)

```
InitialContext ic = new InitialContext();  
WorkManager wm = (WorkManager)ic.lookup("java:comp/env/wm/myWorkManager");
```

The resource reference for the work manager (in this case, wm/myWorkManager) must be declared as a resource reference in the application's deployment descriptor.

4. Call the WorkManager.startWork() method using the work object as a parameter.

For example:

```
Work w = new MyWork(...);  
WorkItem wi = wm.startWork(w);
```



The `startWork()` method can take a `startTimeout` parameter. This specifies a hard time limit for the `Work` object to be started.

The `startWork()` method returns a work item object. This object is a handle that provides a link from the component to the now running work object.

5. [Optional] If your application component needs to wait for one or more of its running work objects to complete, call the `WorkManager.join()` method.

For example:

```
WorkItem wiA = wm.start(workA);
WorkItem wiB = wm.start(workB);
ArrayList l = new ArrayList();
l.add(wiA);
l.add(wiB);
if(wm.join(l, wm.JOIN_AND, 5000)) // block for up to 5 seconds
{

    // both wiA and wiB finished
}
else
{

    // timeout

    // we can check wiA.getStatus or wiB.getStatus to see which, if any, finished.
}
```

This method takes an array list of work items that your component wants to wait on and a flag that indicates whether the component will wait for one or all of the work objects to complete. You also can specify a timeout value.

6. Using the `release()` method, set a variable in a synchronized block.

For example:

```
public synchronized void release()
{
    released = true;
}
```

The `Work.run()` method should periodically examine this variable to check whether the loop should exit or not.

## Work objects

A work object is a type of asynchronous bean used by application components to run code in parallel or in a different J2EE context.

A work object implements the `com.ibm.websphere.asynchBeans.Work` interface. A work object is essentially a `java.lang.Runnable` object that is serializable and provides additional methods. For details, see the Interface `Work` ([../javadoc/ee/com/ibm/websphere/asynchbeans/Work.html](http://../javadoc/ee/com/ibm/websphere/asynchbeans/Work.html)) in the Javadoc.

A component wanting to run work in parallel, or in a different J2EE context, locates a work manager in JNDI, then calls the `WorkManager.startWork()` method using the work object as a parameter.

The `startWork()` method returns a work item object. This object is a handle that provides a link from the component to the now running work object. The work item object is typically used when the component needs to wait for one or more of its running work objects to complete. The `WorkManager.join()` method takes an array list of work items that the component wants to wait on, and a flag indicating

whether the component will wait for all or one of the work objects to complete. A timeout can be specified, which prevents the component from waiting indefinitely.

Why not have the application simply create Java 2 SDK threads? The threads created by the Java 2 SDK are not managed threads and hence know nothing about the J2EE environment and are unusable inside an application server. In addition, these threads have no J2EE context (for example, a `java:comp`) and are not authenticated when they fire. Work object threads, on the other hand, are fully supported by the application server and have the same properties as any other asynchronous bean.

## Example: Work object

The following is an example of a work object that dynamically subscribes to a topic:

```
class SampleWork implements Work
{
    boolean released;
    Topic targetTopic;
    EventSource es;
    TopicConnectionFactory tcf;
    public SampleWork(TopicConnectionFactory tcf, EventSource es, Topic targetTopic)
    {
        released = false;
        this.targetTopic = targetTopic;
        this.es = es;
        this.tcf = tcf;
    }
    synchronized boolean getReleased()
    {
        return released;
    }
    public void run()
    {
        // setup our JMS stuff.
        TopicConnection tc = tcf.createConnection();
        TopicSession sess = tc.createSession(false, Session.AUTOACK);
        tc.start();

        MessageListener proxy = es.getEventTrigger(MessageListener.class, false);
        while(!getReleased())
        {
            // block for up to 5 seconds.
            Message msg = sess.receiveMessage(5000);
            if(msg != null)
            {
                tc.close();
            }
        }
        // called when we want to stop the Work object.
        public synchronized void release()
        {
            released = true;
        }
    }
}
```

As a result, any component that has access to the event source can add an event on demand, which allows components to subscribe to a topic in a more scalable way than by simply giving each client subscriber its own thread. The previous example is fully explored in the WebSphere Trader sample. See your Samples Gallery for details.

---

## Developing event listeners

Application components that listen for events can use the `EventSource.addListener()` method to register an event listener object (a type of asynchronous bean) with the event source to which the events will be published. An event source also can fire events in a type-safe manner using any interface.

Notifications between components within a single EAR file are handled by a special event source. See the topic, [Using the application notification service](#).

### Steps for this task

1. Create an event listener object, which can be any type. For example, see the following interface code:

```
interface SampleEventGroup
{
    void finished(String message);
}

class myListener implements SampleEventGroup
{
    public void finished(String message)
    {
        // This will be called when we 'finish'.
    }
}
```

2. Register the event listener object with the event source

For example, see the following code:

```
InitialContext ic = ...;
EventSource es = (EventSource)ic.lookup("java:comp/websphere/ApplicationNotificationService");
myListener l = new myListener();
es.addListener(l);
```

This enables the `myListener.finished()` method to be called whenever the event is fired. The following code example shows how this event might be fired:

```
InitialContext ic = ...;
EventSource es = (EventSource)ic.lookup("java:comp/websphere/ApplicationNotificationService");
myListener proxy = es.getEventTrigger(myListener.class);
// fire the 'event' by calling the method
// representing the event on the proxy
proxy.finished("done");
```

## Using the application notification service

During an application's lifetime, individual J2EE components (servlets or enterprise beans) within a single EAR file might need to signal each other. There is an event source in the `java:comp` namespace that is bound into all components within an EAR file. The JNDI name for this event source is:

```
java:comp/websphere/ApplicationNotificationService
```

Components within the same application can fire asynchronous events and register event listeners using this application notification service. Startup beans can be used to register these event listeners at application startup or they can be registered dynamically at run time.

To have your enterprise bean or servlet use the application notification service, write code similar to what is shown in the following example:

```
InitialContext ic = new InitialContext();
EventSource appES = (EventSource)
    ic.lookup("java:comp/websphere/ApplicationNotificationService");
// now, the application can add a listener using the EventSource.addListener method.
// MyEventType is an interface.
MyEventType myListener = ...;
AppES.addListener(myListener);

// later another component can fire events as follows
InitialContext ic = new InitialContext();
EventSource appES = (EventSource)
    ic.lookup("java:comp/websphere/ApplicationNotificationService");

// This highlights a constant string on the EventSource interface which
// specifies the 'java:comp/websphere/ApplicationNotificationService' string.
ic.lookup(appES.APPLICATION_NOTIFICATION_EVENT_SOURCE)
// now, the application can add a listener using the EventSource.addListener method.
MyEventType proxy = appES.getEventTrigger(MyEventType.class, false);
proxy.someEvent(someArguments);
```

## Example: Event listener

The following code example demonstrates how to fire a listenerCountChanged event:

```
// imagine this snippet inside an EJB or servlet method.
// Make an inner class implementing the required event interfaces.
EventSourceEvents listener = new Object() implements EventSourceEvents.class
{
    void listenerCountChanged(EventSource es, int old, int newCount)
    {
        try
        {
            InitialContext ic = new InitialContext();
            // here, the asynch bean can access an environment variable of
            // the component which created it.
            int i = (Integer)ic.lookup("java:comp/env/countValue").intValue();
            if(newCount == i)
            {
                // do something interesting
            }
            // this should be called when the code below executes.
        }
        catch(NamingException e)
        {
        }
    }
    void listenerExceptionThrown( EventSource es, Object listener,
        String methodName, Throwable exception)
    {
    }
    void unexpectedException(EventSource es, Object runnable, Throwable exception)
    {
    }
}
// register it.
es.addListener(listener);

...

// now fire an event which the above listener should receive.
EventSourceEvents proxy = (EventSourceEvents)
    es.getEventTrigger(EventSourceEvents.class, false);
```

```
proxy.listenerCountChanged(es, 0, 1);

// now, fire another event, we could call any of the methods.
proxy.listenerCountChanged(es, 4, 5);
```

In this example, we get a proxy for the interface that we want to fire a method on. Then, call the method corresponding to the event on the proxy. This causes the same method, with the same parameters, to be called on any event listeners that implement the `EventSourceEvents` interface and that were previously registered with the `EventSource "es"`. The same proxy can be used to send multiple events simultaneously.

The boolean parameter on the `getEventTrigger()` is named `"sameTransaction"`. When the `sameTransaction` parameter is false, a new transaction is started for each event listener invoked and these event listeners can be called in parallel to the caller. However, the `event()` method always is blocked until all of the event listeners have been notified. If the `sameTransaction` parameter is true, the current transaction, if any, on the thread is used for all of the event listeners; that is, the event listeners share the transaction of the method that fired the event. For that reason, all event listeners must run serially in an undetermined order. That is, the order in which the listeners are called is undefined and the order in which listeners were registered should not be a guide for the order used at run time. The method on the proxy does not return until all of the event listeners have been called; that is, it is a synchronous operation.

The parameters are passed by reference and listeners should not interfere with these references unless the method that fired the event has purposefully designed such interaction. For example, event listeners can be used as collaborators and add data to a map, which was a parameter. Each event listener runs on its own transaction, independent of any transaction that is active on the thread. Extreme care must be taken when the `sameTransaction` parameter is false because the parameters can potentially be accessed by multiple threads.

---

## Developing Asynchronous scopes

Asynchronous scopes are units of scoping that comprise a set of alarms, subsystem monitors, and child asynchronous scopes. Using asynchronous scopes can involve some or all of the following steps:

### Steps for this task

#### 1. Create asynchronous scopes

Create a parent asynchronous scope object by calling the `AsynchScopeManager.createAsynchScope()` method using a unique name as the parameter.

You can store properties in an asynchronous scope object. This provides J2EE applications with a way to store a non-serializable state that otherwise cannot be stored in a session bean.

You also can create child asynchronous scopes, which is useful for scoping data beneath the parent.

#### 2. Listen for alarm notifications

- a. Create a listener object by implementing the `AlarmListener` interface. For more information, see the `AlarmListener` interface (`../javadoc/ee/com/ibm/websphere/asynchbeans/AlarmListener.html`) in the Javadoc.

- b. Supply this object to the `AlarmManager.create()` method, as the target for the alarm.

The `create()` method takes the following parameters:

**Target for the alarm**

The target on which the `fired()` method is called when the alarm is fired.

**Context**

The context object for the alarm. This is useful for supplying alarm-specific data to the listener and allows a single listener to be used for multiple alarms.

**Interval**

The number of milliseconds before the alarm fires.

After the specified interval, the alarm fires and the `fired()` method of the listener is called with the firing alarm as a parameter. The alarm object, itself, is returned. By calling methods on this object, you can cancel or reschedule the alarm.

3. Monitoring remote systems
  - a. Implement a mechanism for detecting messages sent from the remote system. For example, publish-subscribe messaging.
  - b. Create a subsystem manager object by calling the `SubsystemMonitorManager.create()` method with the following parameters:

**Name** Each subsystem monitor must have a unique name.

**Heartbeat interval**

The expected interval, in milliseconds, between heartbeats.

**Missed heart beats until stale or suspect**

The number of heartbeats that can be missed before the subsystem is marked as stale.

**Missed heart beats until dead**

The number of heartbeats that can be missed before the system is marked as dead.

- c. Create an object that implements the `SubsystemMonitorEvents` interface. For more information, see the `SubsystemMonitorEvents` (`../javadoc/ee/com/ibm/websphere/asynchbeans/SubsystemMonitorEvents.html`) in the Javadoc.
- d. Add an instance of this object to the subsystem monitor using the `SubsystemMonitor.addListener()` method.
- e. Whenever a heartbeat message arrives from the remote system, call the `SubsystemMonitor.ping()` method.

The subsystem monitor configures alarms to track the heartbeat status of the remote system. Whenever the `ping()` method is called, the alarms are reset. If an alarm fires, the `ping()` method has not been called; that is, the application did not receive a heartbeat from the subsystem being monitored.

Usage scenario

Asynchronous scopes are useful in stateful server applications. An application can have a startup bean that creates an asynchronous scope on a named work manager. The application also might create subsystem monitors to monitor the health of any remote systems on which the application is dependent.

When a client attaches to the server, the application creates a child asynchronous scope that is owned by the application asynchronous scope for the client and named using the client ID. A subsystem monitor for monitoring the client itself might be created on the client asynchronous scope. If the client times out, a callback can clean up the client state on the server. Callbacks can be attached to the application subsystem monitors, on behalf of the client. When a remote system becomes unavailable, the client code in the server is notified and an event is sent to the client to warn it that a critical remote system has failed. For example, the failure might be a data feed in an electronic trading application.

## Asynchronous scopes

An asynchronous scope (`AsynchScope` object) is a unit of scoping provided for use with asynchronous beans.

Asynchronous scopes are collections of alarms, subsystem monitors, and child asynchronous scopes that enable a relationship to be formed. Each asynchronous scope uses a single work manager.

Each `AsynchScope` object owns and controls the life cycle of the following objects:

### Child asynchronous scopes

Each `AsynchScope` object extends the `AsynchScopeManager` interface, which is a factory for `AsynchScope` objects. (For more information on the `AsynchScopeManager` interface, see the Javadoc (`../javadoc/ee/com/ibm/websphere/asynchbeans/AsynchScopeManager.html`)). Any asynchronous scope can therefore create named asynchronous scopes (children). Child asynchronous scopes can be useful for scoping data underneath the parent. All of the child asynchronous scopes must be uniquely named. These children are destroyed if the parent asynchronous scope is destroyed.

### Alarms

Each asynchronous scope has an associated alarm manager. All of the alarms created by the alarm manager are automatically cancelled if the associated asynchronous scope is destroyed.

### Subsystem monitors

Each asynchronous scope has a subsystem monitor manager, which manages a set of subsystem monitors associated with the asynchronous scope. When the asynchronous scope is destroyed, all of the associated subsystem monitors also are destroyed.

To summarize, asynchronous scopes can be organized into an acyclic tree. The life cycle of each asynchronous scope is directly coupled to that of its parent asynchronous scope. Each asynchronous scope is associated with a set of alarms and subsystem monitors, and an optional set of child asynchronous scopes. These objects are cancelled and destroyed when the asynchronous scope is destroyed.

### Asynchronous scope state

Each asynchronous scope has an associated map, in which applications can store state in the form of name and value pairs.

### Asynchronous scope events

Each asynchronous scope is also an event source. Applications can therefore register event listeners against the asynchronous scope. The event listeners can receive notification if, for example, the `AsynchScope` object is about to be destroyed.

Applications also can use this event source to fire events only to listeners of this asynchronous scope. For example, an `AsynchScope` object created for a client session might be used to fire asynchronous events to parties interested in that client.

## Alarms

An alarm executes J2EE context-aware code at a given time interval. Alarm objects are fine-grained, non-persistent, transient, and can fire at millisecond intervals.

Alarms, themselves, are executed using a thread pool associated with the work manager that owns the associated asynchronous scope.

The `AlarmManager.createAlarm()` method takes an application-written object that implements the `AlarmListener` interface. (For more information on the `AlarmListener` interface, see the Javadoc ([../javadoc/ee/com/ibm/websphere/asynchbeans/AlarmListener.html](http://../javadoc/ee/com/ibm/websphere/asynchbeans/AlarmListener.html).) The fired method is called when the alarm expires. The `createAlarm()` method returns a non-serializable handle, which can be used to cancel or reset the alarm. All of the pending alarms are cancelled when its associated `AsynchScope` object is destroyed.

The Java 2 SDK already has a timer mechanism, so why create a new one? The Java 2 SDK is a J2SE feature that knows nothing about the J2EE environment. Timers fired by the J2SE feature do not run on a managed thread and are therefore unusable inside an application server. These timers also lack a J2EE context (that is, a `java:comp` value) and are not authenticated when they fire. The asynchronous scope alarms are fully supported by WebSphere Application Server Enterprise and have the same properties as any other asynchronous bean.

### Alarm performance

The alarm subsystem is designed to handle a large number of alarms. However, do not have alarms undertake heavy processing when they are firing as this slows the processing of later alarms. If an alarm needs to process a heavy load, design a work object that is activated by a work manager. This procedure moves the heavy processing to a different thread and enables the alarm threads to process alarms unhampered. All of the alarms owned by asynchronous scopes that, in turn, are owned by a single work manager, share a common thread pool. The properties of this thread pool can be tuned at the work manager level using the administrative console.

## Subsystem monitors

A subsystem monitor is an object that monitors the health of a remote system. It uses an event source to inform all registered listeners of the health of the system.

Advanced J2EE applications often rely on remote, non-managed, non-J2EE systems. These remote systems can periodically send clients a message to indicate that they are working. A subsystem monitor is essentially a set of alarms that track indicators messages or "heartbeats" from a remote system.



An application creates a subsystem monitor by calling the `SubsystemMonitorManager.create()` method with the following parameters:

**Name** Each subsystem monitor must be uniquely named.

**Heart beat interval**

The time period, in milliseconds, between arriving heartbeat messages.

**Missed heart beats until stale or suspect**

The number of heartbeats that can be missed before the subsystem is marked as stale. This designation indicates that the subsystem might be having problems.

**Missed heart beats until dead**

The number of heartbeats that can be missed before the system is considered to be down. The system then is marked as dead.

The subsystem monitor configures alarms to track the heartbeat status. Whenever the `ping()` method is called, the alarms are reset. If an alarm fires, the `ping()` method has not been called; that is, the application did not receive a heartbeat from the subsystem being monitored. When the number of **Missed heart beats until stale** has elapsed without a ping, a stale event is fired. Later, if the number of **Missed heart beats until dead** elapses without a ping, a dead event is fired. If a ping is received after a stale or dead notification, a fresh event is sent, which indicates that the subsystem is alive again.

Make **Missed heart beats until dead** greater or equal to the **Missed heart beats until stale**. If **Missed heart beats until stale** equals **Missed heart beats until dead**, a stale event is not published; only a dead event is published.

Applications that want to be informed of these events can register a listener that implements the `SubsystemMonitorEvents` interface. For more information on the `SubsystemMonitorEvents` interface, see the Javadoc ([../javadoc/ee/com/ibm/websphere/asynchbeans/SubsystemMonitorEvents.html](http://../javadoc/ee/com/ibm/websphere/asynchbeans/SubsystemMonitorEvents.html)).

Heart beat messages can be transmitted using a variety of mechanisms. The application must call the `SubsystemMonitor`'s `ping()` method whenever a heartbeat message arrives from a remote system, but the method used to detect these messages is up to the application. For example, you might use a Java Message Service (JMS) publish or subscribe implementation or even a third-party Java messaging product that does not implement JMS.

## Asynchronous scopes: Dynamic message bean scenario

J2EE now supports message-driven beans, but the beans are static. All of the message sources must be known in advance and bound at deployment time. This is not always viable, especially in fluid messaging environments such as those found in brokerages. Some environments have publish-subscribe topic spaces that are continually changing and clients need servers that can subscribe on demand to an arbitrary topic.

An asynchronous bean application can create a work object that performs a blocking receive on a JMS topic and then publishes the message as an event on an application-defined event source. Clients requiring a subscription to that message can add an event listener to the event source. The event source can inform the work object when there are no listeners. Then, the event source can shut down and make the JMS and thread resources available. The work object registers a listener with its own event source. When the count is one again, the work object knows

that it is the only listener and its time to shut down the work object. The WebSphere Trader sample (see your installed Samples Gallery) uses this pattern to dynamically subscribe to JMS topics at run time to gather stock prices. For more information, see an overview of the samples.

How does the server catch clients that disconnect or crash? It creates a subsystem monitor to watch the client and adds an event listener to catch dead events. When a dead event occurs, the server application can clean up the client's server state. For example, the server application can remove the client's event listener from the dynamic message bean; thereby allowing the server to subscribe to a dynamic topic only when it is needed.

---

## Interoperating with asynchronous beans

### Before you begin

The WebSphere Application Server Enterprise Edition Asynchronous Beans 5.0 service has interoperability problems that affect how the dependent services behave in a clustered environment with mixed WebSphere product versions and or mixed platforms. Read the Asynchronous beans interoperability issues article to determine whether you need to complete the following general recovery steps.

### Steps for this task

1. Apply interim fix PQ72742 or PTF 1 to bring all affected servers up to the appropriate Version `WorkWithExecutionContext` level.

At this point, all servers are functioning normally, assuming that the security problems are **not** true.

2. If any of the conditions described in the "Asynchronous beans security errors" section of Asynchronous beans interoperability issues document are true, re-serialize all data. For example, modify the existing EJB or servlet used to serialize the `WorkWithExecutionContext`, so that it rewrites the object using a `java.io.ObjectOutputStream`:

```
ObjectInputStream ois = new ObjectInputStream(...);
ObjectOutputStream oos = new ObjectOutputStream(...);
Object in = ois.readObject();
oos.writeObject(in);
oos.flush();
```

## Asynchronous beans interoperability issues

Asynchronous beans allow taking snapshots of J2EE contextual information from the current servlet or EJB method. This information can include security information, and can be stored using the `WorkManager.create(Work r)` method to get a serializable `WorkWithExecutionContext` object.

If this `WorkWithExecutionContext` object is serialized and stored with customer data, the following information must be reviewed to determine if any steps must be taken to avoid problems.

### **Affected products**

Both the Scheduler and Process Choreographer services that are included with the WebSphere Application Server Enterprise Edition 5 are affected by the problems described in this document. Review the respective sections of each service in the Version 5.0.1 Release Notes or in the one following interim fix options for further information:

- PQ72885 - Scheduler Version 5 Interoperability Issues

- PQ72886 - Process Choreographer Version 5 Interoperability Issues

### **Asynchronous beans security errors**

The asynchronous beans `WorkWithExecutionContext` object created with the 5.0 version of asynchronous beans stores security information incorrectly, which can cause deserialization errors and `ASYN9999E` errors to appear in the WebSphere Application Server log files. This happens if the interoperating versions of WebSphere Application Server have different implementations of Java Authentication and Authorization Service (JAAS). This error most likely occurs when interoperating between two different platforms, but may occur on the same platform if the JAAS implementation is updated on any of the servers.

Security information is stored with the `WorkWithExecutionContext` object only when the following conditions are true:

- The `WorkManager` resource used to create the `WorkWithExecutionContext` object has the *security* context policy attribute enabled.
- Global Security is enabled on the WebSphere Application Server.

If either of the previous conditions are false, then there should be no interoperability issues.

If the JAAS implementation JAR files have already been updated or changed such that all servers are already experiencing `ASYN9999E` errors, then apply the original version of the JAAS implementation JAR files that were used to serialize the `WorkWithExecutionContext` data, or follow the Recovery and interoperability steps in the article, (Interoperating with asynchronous beans), on each platform that is having the problem to ensure all data is reserialized to the appropriate version.

### **Asynchronous beans interoperability errors**

To avoid further problems interoperating between releases or platforms of asynchronous beans, the serialization mechanisms within asynchronous beans have been updated such so that a 5.0 version of the asynchronous beans object, `WorkWithExecutionContext`, is not able to deserialize a `WorkWithExecutionContext` object that is serialized with a Version 5.0.1 release or later.

### **Interoperability and recovery**

If any of the conditions described in the Asynchronous beans security errors section of this document are true, then all data with the serialized Version 5 `WorkWithExecutionContext` object must be reserialized with the updated format.

Regardless of whether any of the conditions described in the "Asynchronous beans security errors" section of this document are true, apply interim fix, PQ72742, to all WebSphere Application Server Enterprise 5 servers that expect to interoperate with Version 5.0.1 and later servers or expect to interoperate with other Version 5 servers that have PQ72742 applied.

Servers with interim fix, PQ72742, or Version 5.0.1 applied are able to read `WorkWithExecutionContext` data serialized with Version 5, but servers at the Version 5 level are not able to read data serialized in the PQ72742 or Version 5.0.1 format. If this problem occurs, errors such as the following appear: `ASYN9999E: Unexpected Exception Occurred: java.io.OptionalDataException`.

For more information on completing the steps for recovery and interoperability, see the article, (Interoperating with asynchronous beans).

## Internationalization interoperability issues

There are internationalization interoperability issues with the Asynchronous Beans `WorkWithExecutionContext` object. Read this section and the Internationalization Interoperability section of the Version 5.0.1 Release Notes for details on how to determine if this issue affects your work environment, and if so, how to resolve the problems.

Affected customers are those who store Internationalization service context information with the `WorkWithExecutionContext` object. Internationalization information is stored with the `WorkWithExecutionContext` object only when the following conditions are true:

- The `WorkManager` resource used to create the `WorkWithExecutionContext` object has the **com.ibm.ws.i18n** context policy attribute enabled.
- The Internationalization service is enabled on the WebSphere Application Server.

If either of the previous conditions are false on either the application server that *serializes* the `WorkWithExecutionContext` or the application server that *deserializes* the `WorkWithExecutionContext`, no interoperability issue exists.

If both of the previous conditions are true on either of the application servers, then follow the instructions in the Internationalization Interoperability section of the Version 5.0.1 Release Notes or interim fix PQ73371.

---

## Chapter 18. Using object pools

An object pool enables an application to avoid creating new Java objects repeatedly. Most objects can be created once, used, and then reused at a later point. An object pool allows an object to be pooled while waiting for the point when it can be reused. These object pools are not meant to be used for pooling JDBC connections or JMS connections and sessions. WebSphere provides specialized mechanisms for dealing with those types of objects. These object pools are intended for pooling application-defined objects or basic JDK types.

To use an object pool, the WebSphere administrator must define an *object pool manager* using the administrative console. Multiple object pool managers can be created in a Websphere cell.

**Note:** The Object pool manager service is only supported from within the EJB Container or Web Container (EJBs or Servlets). Looking-up and using a configured object pool manager from a J2EE application client container is not supported.

### Steps for this task

1. Start the administrative console.
2. Select **Resources > Object Pools**.
3. Define the name of the object pool manager. This name can be up to 30 ASCII characters long.
4. Assign the object pool manager a JNDI name.
5. Provide a description of this object pool manager.
6. Categorize the object pool manager.

### Results

After completing this steps, applications can find the object pool manager by doing a JNDI lookup using the specified JNDI name.

### Usage scenario

The following code illustrates how an *application* can find an Object pool manager object:

```
InitialContext ic = new InitialContext();
ObjectPoolManager opm = (ObjectPoolManager)ic.lookup("java:comp/env/pool");
```

Once the application has an ObjectPoolManager, it can cache an object pool for classes of the types it wants to use. The following is an example:

```
ObjectPool arrayListPool = null;
ObjectPool vectorPool = null;
try
{
    arrayListPool = opm.getPool(ArrayList.class);
    vectorPool = opm.getPool(Vector.class);
}
catch(InstantiationException e)
{
    // problem creating pool
}
```

```

catch(IllegalAccessException e)
{
    // problem creating pool
}

```

Once the application has the pools, it can use them as in the following example:

```

ArrayList list = null;
try
{
    list = (ArrayList)arrayListPool.getObject();
    list.clear(); // just in case
    for(int i = 0; i < 10; ++i)
    {
        list.add("" + i);
    }
    // do what ever we need with the ArrayList
}
finally
{
    if(list != null) arrayListPool.returnObject(list);
}

```

This is the basic pattern for using object pooling. If the application "forgets" to return the object, the only adverse effect is that the object cannot be reused.

---

## Object pool managers

Object pool managers control the reuse of application objects and JDK objects such as Vectors and HashMaps.

Multiple object pool managers can be created in a WebSphere cell. Each object pool manager has a unique cell-wide JNDI name. Applications can find a specific object pool manager by doing a JNDI lookup using the specific JNDI name.

The Object pool manager and its associated objects implement the following interfaces:

```

public interface ObjectPoolManager
{
    ObjectPool getPool(Class aClass)
        throws InstantiationException, IllegalAccessException;
    ObjectPool createFastPool(Class aClass)
        throws InstantiationException, IllegalAccessException;
}

public interface ObjectPool
{
    Object getObject();
    void returnObject(Object o);
}

```

Each object pool manager can be used to pool any Java object with the following characteristics:

- The object must be a public class with a public default constructor.
- Each object class to be pooled must have its own object pool.
- An application gets an object pool for a specific object using either the `ObjectPoolManager.getPool()` or `ObjectPoolManager.createFastPool()` method. The difference between these methods is that the `getPool()` method returns a pool

that can be shared across multiple threads. The `createFastPool()` method returns a pool that can only be used by a single thread.

If, in a JVM, the `getPool()` method is called multiple times for a single class, the same pool is returned. A new pool is returned for each call when the `createFastPool()` method is called. Basically, the `getPool()` method returns a pool that is thread-synchronized.

The pool for use by multiple threads is slightly slower than a fast pool due to the need to handle thread synchronization. However, extreme care must be taken when using a fast pool. Consider the following interface:

```
public interface PoolableObject
{
    void init();
    void returned();
}
```

If the objects placed in the pool implement this interface and the `ObjectPool.getObject()` is called, the object returned has the `init()` method called on it. When the `ObjectPool.returnObject()` method is called, the `returned()` method is called on the object before it is returned to the object pool. This allows objects to be pre-initialized or cleaned up.

It is not always possible for an object to implement `PoolableObject`. For example, an application might want to pool `ArrayList` objects. The `ArrayList` would need to be cleared each time the application reuses it. The application might extend `ArrayList` and have that implement `Poolable`. For example, consider the following:

```
public class PooledArrayList extends ArrayList implements PoolableObject
{
    public PooledArrayList()
    {
    }

    public void init() {
    }

    public void returned()
    {
        clear();
    }
}
```

If the application uses this, in place of a true `ArrayList`, the `ArrayList` is cleared automatically when it is returned to the pool.

**Note:** Clearing an `ArrayList` simply marks it as empty and the array backing the `ArrayList` is not freed.

Therefore, as the application reuses the `ArrayList`, the backing array expands until it is big enough for all of the application requirements. Once this point is reached, it stops allocating and copying new backing arrays and achieves the best performance.

It might not be possible or desirable to use the previous procedure. An alternative is to implement a custom object pool and register this with the object pool manager as the pool to use for classes of that type. The class is registered by the WebSphere administrator when the object pool manager is defined in the cell. Take care that these classes are packaged in JAR files available on all of the nodes in the cell where they might be used.

---

## Object pool manager collection

Use this page to manage object pool managers.

To view this administrative console page, click **Resources > Object Pools** .

### Name

The name by which the object pool manager is known for administrative purposes.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Range     | 1 through 30 ASCII characters |

### JNDI Name

The JNDI name for the object pool manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

### Description

A description of the object pool manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

### Category

A category string used to classify or group this object pool manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

## Object pool manager settings

Use this page to modify object pool manager settings.

To view this administrative console page, click **Resources > Object Pools > *objectpoolmanager\_name***

### Name

The name by which the object pool manager is known for administrative purposes.

|           |                               |
|-----------|-------------------------------|
| Data type | String                        |
| Range     | 1 through 30 ASCII characters |

### JNDI Name

The JNDI name for the object pool manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

### Description

A description of the object pool manager.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|



## Category

A category string used to classify or group this object pool manager.

Data type String

## Custom object pool collection

Use this page to manage object pools.

To view this administrative console page, click **Resources > Object Pools > *objectpoolmanager\_name* > Object Pools**.

**Pool Class Name:** The fully-qualified class name of the objects that are stored in the object pool.

Data type String

**Pool Impl Class Name:** The fully-qualified class name of the CustomObjectPool implementation class for this object pool.

Data type String

## Custom object pool settings

Use this page to modify custom object pool settings.

An object pool manages a pool of arbitrary objects.

To view this administrative console page, click **Resources > Object Pools > *objectpoolmanager\_name* > Object Pools > *objectpool\_name***.

**Pool Class Name:** The fully-qualified class name of the objects that are stored in the object pool.

Data type String

**Pool Impl Class Name:** The fully-qualified class name of the CustomObjectPool implementation class for this object pool.

Data type String

---

## Object pool service settings

Use this page to enable or disable the object pool service, which manages object pool resources used by the server.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Object Pool Service** .

## Startup

Specifies whether the server will attempt to start the object pool service.

Default Selected

Range

**Selected**

When the application server starts, it attempts to start the object pool service automatically.

**Cleared**

The server does not try to start the object pool service service. If object pool resources are to be used on this server, the system administrator must start the object pool service manually or select this property then restart the server.

---

## Object pools: Resources for learning

Use the following links to find relevant supplemental information about object pools. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

### Programming model and decisions

•



**Java theory and practice: Thread pools and work queues** (<http://www-106.ibm.com/developerworks/library/j-jtp0730.html>)

•



**Java performance programming, Part 1: Smart object-management saves the day** (<http://www-106.ibm.com/developerworks/library/jw-performance.html>)

•



**Build your own ObjectPool in Java to boost app speed**  
(<http://www.javaworld.com/jw-06-1998/jw-06-object-pool.html>)

•



**Improve the robustness and performance of your ObjectPool**  
(<http://www.javaworld.com/jw-08-1998/jw-08-object-pool.html>)

•



**Java Tip 78: Recycle broken objects in resource pools**  
(<http://www.javaworld.com/javaworld/jvatips/jw-jvatip78.html>)

---

## Object pool performance considerations

Object pools are designed to pool and reuse instances of objects that are both complex and frequently instantiated during runtime. To use an object pool, you must define an object pool manager using the administrative console. The object pool manager is available in the Java Naming Directory Interface (JNDI) and able to provide object pool instances for any class implementing the `PoolableObject` interface. The `PoolableObject` interface defines methods for releasing resources when an object is to be pooled and initializes the internal state when an object is retrieved from the pool. The objects are obtained by methods available on the `ObjectPool` object.

You should use an object pool when objects are needed repeatedly that are expensive to instantiate. This expense can be gauged through the complexity of the object, as well as the size of the data inside the object. Another consideration is the viability of efficiently cleaning out an object. If the members of an object are large or complex, but cannot be quickly restored to defaults for future use, pooling these objects proves effective. The key to seeing a benefit from using an object pool is to make sure that the total time to retrieve the object from the pool, initialize the internal members, and clear the object for reuse and return it to the pool, is less than the time it takes to instantiate the object, initialize internal members, and subsequently garbage collect it when heap space runs out. Retrieving an object from the pool and returning it to the pool have been streamlined as much as possible, but the tasks of member initialization and clearing must fall to the developer to implement and evaluate.

An important issue to consider is when implementing an object pool thread contention. There are two types of object pools to choose from; one is thread-safe and other is not. The latter is not synchronized, allowing for faster access and reduction in contention, but is not recommended in a multi-threaded environment. If there is a chance that two threads are operating on a pool at the same time, use the thread-safe object pool. However, be sure to consider how many threads are accessing the pool at the same time and how frequently. If it is more than a few on a regular basis, the single point of contention created in this scenario can sometimes cancel out all benefits with an otherwise well-conceived poolable object. A thread-safe pool accessed by multiple threads creates a bigger pool of ready object instances, but lining up to retrieve these instances through synchronization can degrade performance.

When attempting to maximize performance with an object pool, cache the `ObjectPoolManager` instance wherever it is accessed. By looking it up once and keeping a reference to it for the future, you can reduce the amount of time spent looking it up in JNDI, which, because of its remote nature, tends to be quite costly.

Another performance consideration is in the implementation of a poolable object's `returned()` method. Preserve the member to be cleared, if it is needed for next use. This reduces even more trash for the garbage collector and avoids the re-initialization of this member upon reuse. Also, if a member has a more efficient way to clear its state, other than simply destroying it, try to utilize the poolable object `returned()` method. For example, calling the `clear()` method on a member `Vector` is much more efficient than destroying the `Vector` and instantiating a new one upon reuse.

Using an object pool significantly reduces the amount of memory that applications use and the amount of garbage created. Through an intelligent choice of pooled

objects and pool types, performance improvements can be achieved by avoiding unnecessary object allocation and garbage collection.

---

## Chapter 19. Using startup beans

A startup bean is a stateful session bean that is loaded when an application starts. Startup beans enable J2EE applications to execute business logic automatically, whenever an application starts or stops normally.

Startup beans are especially useful when used in combination with asynchronous bean features. For example, a startup bean might create an alarm object that uses JMS to periodically publish heartbeat messages on a well-known topic. This enables clients or other server applications to determine whether the application is available.

### Steps for this task

1. Use the home interface, `com.ibm.websphere.startupservice.AppStartupHome`, to designate a bean as a startup bean
2. Use the remote interface, `com.ibm.websphere.startupservice.AppStartup`, to define a `start()` and `stop()` method on the bean.

The bean's `start()` method is called when the application starts. It implements any business logic that needs to run at application start time.

The `start()` method returns a boolean. True indicates normal application startup and false indicates that the application start process should be aborted. The `start()` and `stop()` methods should not use a `TX_MANDATORY` attribute because there never is a transaction on the thread when the `start()` or `stop()` methods are invoked. Any other `TX_*` attribute can be used. If `TX_MANDATORY` is used, an exception is logged (need a transaction for mandatory) and the application does not start.

The bean's `stop()` method is called when the application stops and implements any business logic that needs to run at this time. Any exception thrown by a `stop()` method is ignored, but logged to trace.

The `start()` and `stop()` methods on the remote interface use **Run-As** mode. **Run-As** mode specifies the credential information to be used by the security service to determine the permissions that a principal has on various resources. If security is on, the **Run-As** mode needs to be defined on all of the methods called. The identity of the bean without this setting is undefined. For more information about the **Run-As** mode, see the topic Method extension assembly settings.

There are no restrictions on what code the `start()` and `stop()` methods can run, since the full Enterprise Application Server programming model is available to these methods.

3. Use an *optional* environment property called `wasStartupPriority`, which is an integer, to specify the start order of multiple startup beans in the same JAR file. If the environment property is found and is the wrong type, application startup is aborted. If no priority value is specified, a default priority of 0 is used. It is recommended that you specify the priority property. Beans that have specified a priority are sorted using this property. Beans with numerically lower priorities are executed first. Beans that have the same priority are executed in an undefined order. All priorities must be positive integers. The priority is used to order beans within an EJB JAR file. The order in which this process is applied to different EJB JAR files in a single EAR file is undefined. Beans are stopped in the opposite order to their start priority.

Startup beans must specify a timeout value of 0. Failure to do so causes the bean to be passivated and results in errors when attempting to call the stop() method when the application is stopped.

---

## Chapter 20. Using the scheduler service

The scheduler service enables tasks to be executed at a requested time. The following tasks can be scheduled:

- Invoke a session bean method.
- Send a JMS (Java Message Service) message on a Queue or Topic.

The scheduler service performs the task, repeating as necessary, according to the meta data for that task.

### Steps for this task

1. Developing and scheduling tasks.

This article includes instructions for developing various types of tasks, receiving notifications from a scheduler, submitting tasks to a scheduler, and managing tasks.

**Note:** Creating and manipulating scheduled tasks through the Scheduler interface is only supported from within the EJB Container or Web Container (enterprise beans or servlets). Looking-up and using a configured scheduler from a J2EE application client container is not supported.

2. Managing the scheduler service.

This article includes instructions for creating and configuring a database for scheduler, configuring a scheduler instance, and enabling or disabling the scheduler service (the service is enabled by default).

3. Interoperating with the Scheduler service.

This article explains how to manage the Scheduler service in a clustered environment with mixed WebSphere Application Server product versions and or mixed platforms.

---

## Managing the scheduler service

Schedulers are configured using the administrative console. Schedulers are available to all servers on which the scheduler service is enabled.

### Steps for this task

1. "Creating the database for scheduler".
2. "Configuring a scheduler".
3. "Enabling the scheduler service".

The scheduler service is enabled by default.

## Creating the database for scheduler

### Before you begin

Your database system must be installed and available.

It is important to realize that the scheduler uses this database for storing tasks and then executing them. The performance of the scheduler is ultimately limited by the performance of the database. If you need more tasks per second, you can run the scheduler daemons on larger systems or you can use clusters for the session beans

used by the tasks. Eventually, however, the task database becomes saturated and you then need a larger or better-tuned database system.

Multiple applications can share a scheduler database. This can lower the cost of administering the scheduler database.

Scheduler requires a database, JDBC provider, and data source.

#### Steps for this task

1. Create the database according to the description for your database system:
  - Creating a Cloudscape database for scheduler.
  - Creating a DB2 database for scheduler.
  - Creating an Informix database for scheduler.
  - Creating a Microsoft SQL Server database for scheduler.
  - Creating an Oracle database for scheduler.
  - Creating a Sybase 12.0 database for scheduler.
2. If the database is not on the same machine as your IBM WebSphere Application Server, verify that you can access the database from your application server machine.
3. Configure your JDBC provider and data source.  
For details, see the topic [Creating and configuring a JDBC provider and data source](#).

### **Creating a Cloudscape database for scheduler**

Cloudscape is a database system implemented in Java. It is delivered with WebSphere Application Server as three JAR files. The Cloudscape license that comes with WebSphere is only for development and test, not for production purposes.

#### Steps for this task

1. Open a command-line window.
2. Make sure that you have administrator rights for the database system.
3. If you want to use an existing database, skip to step 5.  
**Note:** Make sure that the database supports Unicode (UTF-8) . Otherwise, it cannot store all characters that can be handled in Java, and you could run into codepage conversion problems when a client uses an incompatible codepage.
4. Use the Cloudview utility supplied with Cloudscape to create a database named scheddb.  
**Note:** Cloudscape allows only one local connection. If WebSphere Application Server is running and accessing a Cloudscape database, attempts to open a second connection to the database from the command line are rejected.
5. Create the schema.
  - a. Using a text editor, edit the script  
`%WAS_HOME%\Scheduler\createSchemaCloudscape.ddl` according to the instruction at the top of the file.
  - b. Enter one of the following commands (shown here on multiple lines for publication):  
On Windows:  

```
%WAS_HOME%\java\jre\bin\java -Djava.ext.dirs=%WAS_HOME%/lib  
-Dij.protocol=jdbc:db2j: -Dij.database=scheddb com.ibm.db2j.tools.ij  
%WAS_HOME%\Scheduler\createSchemaCloudscape.ddl
```



On UNIX:

```
%WAS_HOME%/java/jre/bin/java -Djava.ext.dirs=%WAS_HOME%/lib
-Dij.protocol=jdbc:db2j: -Dij.database=scheddb com.ibm.db2j.tools.ij
%WAS_HOME%/Scheduler/createSchemaCloudscape.ddl
```

## Results

The Cloudscape database for Scheduler exists.

### **Creating a DB2 database for scheduler**

#### Steps for this task

1. Open a DB2 command-line window.
2. Make sure that you have administrator rights for the database system.
3. If you want to use an existing database, skip to step 5.

**Note:** Make sure that the database supports Unicode (UTF-8) . Otherwise, it cannot store all characters that can be handled in Java, and you could run into codepage conversion problems when a client uses an incompatible codepage.

To avoid deadlocks, be sure that the DB2 isolation level is set to "read stability". If necessary, enter the command `db2set DB2_RR_TO_RS=YES` then restart the DB2 instance to activate the change.

4. In the DB2 command line processor, enter this command to create the database:  
`db2 CREATE DATABASE scheddb USING CODESET UTF-8 TERRITORY en-us`

A DB2 database named scheddb has been created.

5. Create the tablespace and schema.
  - a. **(Optional)** Analyze the results of your experiences during development and system testing.

The size of your database depends on many factors. If possible, distribute tablespace containers across different logical disks, and implement an appropriate security policy. Consider the performance implications of your choices for bufferpools and log file settings.

- b. Using a text editor, edit the following scripts according to the instruction at the top of each file: `%WAS_HOME%\Scheduler\createTablespaceDB2.ddl`, `%WAS_HOME%\Scheduler\createSchemaDB2.ddl`, `%WAS_HOME%\Scheduler\dropSchemaDB2.ddl`, and `%WAS_HOME%\Scheduler\dropTablespaceDB2.ddl`.

- c. Make sure that you are attached to the correct instance.

Check the environment variable `DB2INSTANCE`.

- d. To connect to a database named scheddb, enter the command:

```
db2 connect to scheddb
```

- e. To create the tablespace, enter the command:

```
db2 -tf createTablespaceDB2.ddl
```

Make sure that the script's output contains no errors. If there were any errors, you can drop the tablespace using the script `dropTablespaceDB2.ddl`.

- f. To create the schema (tables and indices), in the DB2 command line processor, enter the command:

```
db2 -tf createSchemaDB2.ddl
```

Make sure that the script's output contains no errors. If there were any errors, you can use `dropSchemaDB2.ddl` to drop the schema.

## Results

The DB2 database for scheduler exists.

### **Creating an Informix database for scheduler**

#### Steps for this task

1. Open a command-line window.
2. Make sure that you have administrator rights for the database system.
3. If you want to use an existing database, skip to step 5.  
**Note:** Make sure that the database supports Unicode (UTF-8) . Otherwise, it cannot store all characters that can be handled in Java, and you could run into codepage conversion problems when a client uses an incompatible codepage.
4. If you want to create a new database named `scheddb`, enter the command:

```
dbaccess CREATE DATABASE scheddb with log
```

5. Create the schema.
  - a. Using a text editor, edit the script `%WAS_HOME%\Scheduler\createSchemaInformix.sql` according to the instruction at the top of the file.
  - b. Enter the command:

```
dbaccess scheddb createSchemaInformix.sql
```

## Results

The Informix database for scheduler exists.

### **Creating a Microsoft SQL Server database for scheduler**

#### Steps for this task

1. Open a command-line window.
2. Change to the directory where the configuration scripts for scheduler are located. This is the Scheduler subdirectory of the IBM WebSphere Application Server installation directory.

On Windows, enter:

```
cd %WAS_HOME%\Scheduler
```

On UNIX, enter:

```
cd $WAS_HOME/Scheduler
```

3. Using a text editor, edit the schema creation script (`createSchemaMSSQL.sql` for SQL Server 2000 and `createSchemaMSSQL7.sql` for SQL Server 7), according to the instructions at the beginning of the file.
4. If you want to use an existing database, skip to step 6.  
**Note:** Make sure that the database supports Unicode (UTF-8) . Otherwise, it cannot store all characters that can be handled in Java, and you could run into codepage conversion problems when a client uses an incompatible codepage.
5. If you want to create a new database named `scheddb`:
  - a. Make sure that you are using a user ID that has administrator rights for the database system.

- b. In the **Enterprise Manager**, expand a server group, then expand a server.
- c. Right-click **Databases**, then click **New Database**.
- d. Type the name scheddb.
- e. Modify any default values, as desired, then save.

An Microsoft SQL Server database named scheddb has been created.

6. To create the schema:
  - a. Make sure that you have administrator rights for the database system.  
The user ID you use to create the schema must be the one that you tell WebSphere to use when accessing the database.
  - b. Run the script to create the schema (tables and views):

```
isql -S <serverName> -U<userid> -P<password> -D<databaseName> -i <script name>
```

### Results

The Microsoft SQL Server database for scheduler exists.

## **Creating an Oracle database for scheduler**

### Steps for this task

1. Open a command-line window.
2. Make sure that you have administrator rights for the database system.
3. If you want to use an existing database, skip to step 5.  
**Note:** Make sure that the database supports Unicode (UTF-8) . Otherwise, it cannot store all characters that can be handled in Java, and you could run into codepage conversion problems when a client uses an incompatible codepage.
4. Use the Database Configuration Assistant to create a database named scheddb.  
Make sure that you select the JServer option for the database. It is recommended to use a Unicode codepage when creating the database. The text data you pass to the APIs must be compatible with the selected codepage.
5. Create the tablespace and schema.
  - a. Using a text editor, edit the scripts  
%WAS\_HOME%\Scheduler\createTablespaceOracle.ddl and  
%WAS\_HOME%\Scheduler\createSchemaOracle.ddl according to the instruction at the top of the files.
  - b. If you do not want the schema to be created in the default instance, set the environment variable ORACLE\_SID
  - c. To create the tablespace, run the script createTablespaceOracle.ddl.  
For test purposes you can use the same location for all tablespaces and pass the path as a command line argument to the script, for example, on Windows, user ID scheduser, password schedpwd, database name scheddb, and tablespace path d:\mydb\ts, enter the command:  
sqlplus scheduser/schedpwd@scheddb @createTablespaceOracle.ddl d:\mydb\ts  
  
If you get any errors creating the tablespace, you can use dropTablespaceOracle.ddl to drop the tablespace.
  - d. To create the schema, run the script createSchemaOracle.ddl.  
For example, on Windows, enter:  
sqlplus scheduser/schedpwd@scheddb @createSchemaOracle.ddl

If you get any errors creating the schema (tables and views), you can use dropSchemaOracle.ddl to drop the schema.

## Results

The Oracle database for scheduler exists.

### **Creating a Sybase database for scheduler**

#### Steps for this task

1. Open a command-line window.
2. Make sure that you have administrator rights for the database system.
3. Make sure that you have the DTM option for Sybase ASE installed.
4. If you want to use an existing database, skip to step 6.

**Note:** Make sure that the database supports Unicode (UTF-8) . Otherwise, it cannot store all characters that can be handled in Java, and you could run into codepage conversion problems when a client uses an incompatible codepage.

5. Use the Sybase isql utility to create a database named scheddb. See your Sybase product documentation for details.
6. Create the schema:
  - a. Using a text editor, edit the script  
%WAS\_HOME%\Scheduler\createSchemaSybase12.ddl according to the instruction at the top of the file.
  - b. Enter the command:

```
isql -S <serverName> -U<userid> -P<password> -D scheddb -i createSchemaSybase12.ddl
```

## Results

The Sybase database for scheduler exists.

## **Configuring a scheduler**

Before your application can make use of the scheduler service, you need to configure a scheduler instance using the administrative console. Conceptually, a scheduler is similar to a datasource: you specify various configuration attributes, including a JNDI name where the instance will be bound. Once defined, an application using the scheduler API can look up the scheduler object and call various methods to manage tasks.

#### Steps for this task

1. Start the administrative console.
2. Select **Resources > Scheduler Configurations**.
3. Click **New**.
4. Specify configuration settings.

Fields marked with an asterisk (\*) are required. The settings are described in detail in the topic Scheduler configuration settings.

### **Scheduler configuration collection**

Use this page to manage scheduler configurations.

To view this administrative console page, click **Resources > Scheduler Configurations** .

**Name:** The name by which this scheduler is known for administrative purposes.

Data type

String

**JNDI Name:** The JNDI name for the scheduler.

The JNDI name specifies where this scheduler instance is bound in the namespace. Clients can look this name up directly, although the use of resource references is recommended.

Data type String

**Description:** A description of this scheduler for administrative purposes.

Data type String

**Category:** A string that can be used to classify or group this scheduler.

Data type String

**Datasource JNDI Name:** Datasource where persistent tasks will be stored.

Any datasource available in the name space can be used with a scheduler. Multiple schedulers can share a single datasource while using different tables by specifying a table prefix.

Data type String

**Datasource Alias:** Alias to a user name and password used to access the datasource.

Data type String

**Table Prefix:** String prepended to the table name TASK.

Multiple independent schedulers can share the same database if each instance specifies a different prefix string.

Data type String

**Poll Interval:** The interval at which the scheduler daemon polls the database. Each scheduled, repeating task's minimum repeat interval will be equal to this value regardless of what is specified on the task.

Each poll operation can be expensive. If the interval is extremely small and there are many scheduled tasks, polling can consume a large portion of system resources.

|           |                           |
|-----------|---------------------------|
| Data type | Integer                   |
| Units     | Seconds                   |
| Default   | 30                        |
| Range     | Any positive long integer |

**Work Manager:** Specifies the work manager used by this scheduler.

The Work Manager is a server object that serves as a logical thread pool for the scheduler. Each repeating task that is created using this scheduler will use the "Number Of Alarm Threads" specified in the Work Manager which will affect the number tasks that can run concurrently. Use the Work Manager's "Service Names" property to limit the amount of context information that is propagated to the task when it executes.

When a task fires, the task is run in the Work Manager associated with the scheduler instance. Configuring a scheduler with a specific Work Manager enables you to control how many tasks are actively running at a given time.

### Scheduler configuration settings

Use this page to modify scheduler settings.

To view this administrative console page, click **Resources > Scheduler Configurations > *scheduler\_name***.

**Name:** The name by which this scheduler is known for administrative purposes.

Data type String

**JNDI Name:** The JNDI name for the scheduler.

The JNDI name specifies where this scheduler instance is bound in the namespace. Clients can look this name up directly, although the use of resource references is recommended.

Data type String

**Description:** A description of this scheduler for administrative purposes.

Data type String

**Category:** A string that can be used to classify or group this scheduler.

Data type String

**Datasource JNDI Name:** Datasource where persistent tasks will be stored.

Any datasource available in the name space can be used with a scheduler. Multiple schedulers can share a single datasource while using different tables by specifying a table prefix.

Data type String

**Datasource Alias:** Alias to a user name and password used to access the datasource.

Data type String

**Table Prefix:** String prepended to the table name TASK.

Multiple independent schedulers can share the same database if each instance specifies a different prefix string.

Data type String

**Poll Interval:** The interval at which the scheduler daemon polls the database. Each scheduled, repeating task's minimum repeat interval will be equal to this value regardless of what is specified on the task.

Each poll operation can be expensive. If the interval is extremely small and there are many scheduled tasks, polling can consume a large portion of system resources.

|           |                           |
|-----------|---------------------------|
| Data type | Integer                   |
| Units     | Seconds                   |
| Default   | 30                        |
| Range     | Any positive long integer |

**Work Manager:** Specifies the work manager used by this scheduler.

The Work Manager is a server object that serves as a logical thread pool for the scheduler. Each repeating task that is created using this scheduler will use the "Number Of Alarm Threads" specified in the Work Manager which will affect the number tasks that can run concurrently. Use the Work Manager's "Service Names" property to limit the amount of context information that is propagated to the task when it executes.

When a task fires, the task is run in the Work Manager associated with the scheduler instance. Configuring a scheduler with a specific Work Manager enables you to control how many tasks are actively running at a given time.

### Creating a scheduler resource reference

When a scheduler has been defined in the server configuration, the object instance is bound into the global name space under the configured JNDI name. A resource reference can be used to avoid hardcoding this JNDI name into your application.

You can alternatively create a scheduler resource reference by editing the XML directly. A scheduler resource reference is a J2EE compliant resource that uses the class `com.ibm.websphere.scheduler.Scheduler` as the object type. For information regarding the XML file format, see the J2EE Specification.

#### Steps for this task

1. Start the Application Assembly Tool.
2. Select your application.
3. In the left-hand panel, right-click on **Resource References** and select **New**.
4. On the **General** tab, complete the following fields:

|             |   |
|-------------|---|
| <b>Name</b> | Specify the name suffix. For example, if the scheduler name is <i>MyScheduler</i> , the reference JNDI name is <b>java:comp/env/MyScheduler</b> |
| <b>Type</b> | From the drop-down list select <b>com.ibm.websphere.scheduler.Scheduler</b>   |
5. (Optional) At this time you can also specify the global JNDI name to which this resource reference is bound by entering the JNDI name on the **Bindings** tab.

## Scheduler daemon

A scheduler daemon is a background thread that searches for events in the persistent store.

A scheduler daemon is started for each scheduler defined on each server. If "Scheduler 1" is configured on server1, then there will only be one scheduler daemon running on server1 unless it is cloned. If "Scheduler 1" is defined at the node scope level, then the scheduler will run on each server within that node.

The poll interval determines the frequency at which the persistent store is queried. By default, this value is set to 30 seconds. When a task is found that is scheduled to fire within the current poll interval, an alarm is set. The task then runs as close to this time as possible using an alarm thread from the scheduler's associated work manager. Thus, the number of alarm threads configured on the work manager determines how many concurrent tasks are executed. No tasks are lost. If we reach this limit, then new tasks are simply queued to be executed when an alarm thread becomes available. The actual firing time is dictated by server load and availability of free threads in the alarm thread pool of the associated work manager.

### Scheduler daemons in a cluster

When multiple scheduler daemons are configured to the same table (as is the case in a clustered environment), any of the daemons can find a task and set the timer in its Java Virtual Machine (JVM). The task is executed in the virtual machine where the timer first fires.

## Enabling the scheduler service

### Before you begin

Before an application can make use of the scheduler service, you need to configure a scheduler.

The scheduler service manages all schedulers used by a given server. The scheduler service can be enabled and disabled on a server-by-server basis using the administrative console. The service is enabled by default. If you disable the service on a server, all schedulers configured on that server are no longer available. All lookups fail and all scheduler daemons are inactive.

### Steps for this task

1. Start the administrative console.
2. Select **Servers** > *server\_name* > **Scheduler Service**.
3. Select or clear the **Startup** checkbox in order to enable or disable the service.
4. Click **Save** on the menu bar to save your configuration.

### Results

The change takes effect on the next server restart.

### Scheduler service settings

Use this page to enable or disable the scheduler service, which manages scheduler resources used by the server.

To view this administrative console page, click **Servers** > **Application Servers** > *server\_name* > **Scheduler Service** .



**Startup:** Specifies whether the server will attempt to start the scheduler service.

Default  
Range

Selected

**Selected**

When the application server starts, it attempts to start the scheduler service automatically.

**Cleared**

The server does not try to start the scheduler service. If scheduler resources are to be used on this server, the system administrator must start the scheduler service manually or select this property then restart the server.

---

## Developing and scheduling tasks

### Steps for this task

#### 1. Developing a task.

The scheduler API supports different implementations of the `TaskInfo` interface, each of which can be used to schedule a particular type of work. Refer to one of the following topics for details:

- Developing a task that calls a session bean.
- Developing a task that sends a JMS message. This task object can send a JMS message to either a queue or a topic.

**Note:** Creating and manipulating scheduled tasks through the Scheduler interface is only supported from within the EJB Container or Web Container (enterprise beans or servlets). Looking-up and using a configured scheduler from a J2EE application client container is not supported.

#### 2. Receiving scheduler notifications.

A notification sink is set on a task in order to receive the notification events that are generated by a scheduler when it performs an operation on the task.

#### 3. Submitting a task to a scheduler.

After a `TaskInfo` object has been created, it can be submitted to the scheduler for task creation by calling the `Scheduler.create()` method.

#### 4. Managing tasks with a scheduler.

## Developing a task that calls a session bean

The scheduler API supports different implementations of the `TaskInfo` interface, each of which can be used to schedule a particular type of work. This topic describes how to call a method on a task handler session bean using the `BeanTaskInfo` implementation.

### Steps for this task

#### 1. Create a stateless session bean that implements the `process()` method in the `com.ibm.websphere.scheduler.TaskHandler` remote interface. The `process()` method is called when the task fires.

The Home and Remote interfaces must be set as follows in the bean's deployment descriptor:

- `com.ibm.websphere.scheduler.TaskHandlerHome`

- com.ibm.websphere.scheduler.TaskHandler

2. Create an instance of the BeanTaskInfo interface by using the following Scheduler.createBeanTaskInfo() method:

```
//lookup the scheduler instance to be used
Scheduler scheduler =
    (Scheduler)new InitialContext.lookup("java:comp/env/Scheduler");

BeanTaskInfo taskInfo = scheduler.createBeanTaskInfo();
```

**Note:** Creating a BeanTaskInfo object does not add the task to the persistent store. Rather, it creates a placeholder for the necessary data. The task is not added to the persistent store until the create() method is called on a Scheduler instance, as described in the topic Submitting a task to a scheduler.

3. Set parameters on the BeanTaskInfo object. These parameters define which task is to run, which session bean is called, and so on.

The TaskInfo interface contains various set() methods that you can use to control execution of the task, including when the task will fire and what work the task will do when it fires. For example:

```
//create a date object which represents 30 seconds from now
java.util.Date startDate = new java.util.Date(System.currentTimeMillis()+30000);

//find the session bean to be called when the task executes
Object o = new InitialContext().lookup("java:comp/env/ejb/MyTaskHandlerHome");
TaskHandlerHome home =
    (TaskHandlerHome)javax.rmi.PortableRemoteObject.narrow(o,TaskHandlerHome.class);

//now set the start time and task handler to be called in the task info
taskInfo.setTaskHandler(home);
taskInfo.setStartTime(startDate);
```

The TaskInfo interface ([../javadoc/ee/com/ibm/websphere/scheduler/TaskInfo.html](http://../javadoc/ee/com/ibm/websphere/scheduler/TaskInfo.html)) specifies additional control points, as documented in Javadoc.

### Results

A TaskInfo object has been created that contains all of the relevant data for a task.

### What to do next

Submit the task to a scheduler instance for creation, as described in the topic Submitting a task to a scheduler.

## Developing a task that sends a JMS message

The scheduler API supports different implementations of the TaskInfo interface, each of which can be used to schedule a particular type of work. This topic describes how to use the MessageTaskInfo implementation, which sends a JMS message to either a queue or a topic.

### Steps for this task

1. Create an instance of the MessageTaskInfo interface by using the following Scheduler.createMessageTaskInfo() method:

```
//lookup the scheduler instance to be used
Scheduler scheduler =
    (Scheduler)new InitialContext.lookup("java:comp/env/Scheduler");

MessageTaskInfo taskInfo = scheduler.createMessageTaskInfo();
```

**Note:** Creating a `MessageTaskInfo` object does not add the task to the persistent store. Rather, it creates a placeholder for the necessary data. The task is not added to the persistent store until the `create()` method is called on a `Scheduler` instance, as described in the topic [Submitting a task to a scheduler](#).

2. Set parameters on the `MessageTaskInfo` object.

The `TaskInfo` interface contains various `set()` methods that can be used to control execution of the task, including when the task will fire and what work the task will do when it fires. For example:

```
//create a date object which represents 30 seconds from now
java.util.Date startDate = new java.util.Date(System.currentTimeMillis()+30000);
```

```
//now set the start time and the JNDI names for the queue connection
// factory and the queue
taskInfo.setConnectionFactoryJndiName("jms/MyQueueConnectionFactory");
taskInfo.setDestination("jms/MyQueue");
taskInfo.setStartTime(startDate);
```

The `TaskInfo` interface ([../javadoc/ee/com/ibm/websphere/scheduler/TaskInfo.html](http://../javadoc/ee/com/ibm/websphere/scheduler/TaskInfo.html)) specifies additional control points, as documented in Javadoc.

### Results

A `TaskInfo` object has been created that contains all of the relevant data for a task.

### What to do next

Submit the task to a scheduler instance for creation, as described in the topic [Submitting a task to a scheduler](#).

## Receiving scheduler notifications

Various notification events are generated by a scheduler when it performs an operation on a task. These events include:

### **Scheduled**

A task has been scheduled.

### **Purged**

A task has been permanently deleted from the persistent store.

### **Suspended**

A task was suspended.

### **Resumed**

A task was resumed.

### **Complete**

A task has run completely. If it was a repeating task, all repeats have been performed.

### **Cancelled**

A task has been cancelled. It will not run again.

**Fired** A task fired successfully.

### **Fire Failed**

A task could not fire successfully.

To receive notification events, call the `setNotificationSink()` method on the `TaskInfo` interface before creating the event. The `setNotificationSink()` method enables you to specify the session bean that is to act as the callback, and a mask that restricts which events are generated.

#### Steps for this task

1. Create a notification sink session bean.

Create a stateless session bean that implements the `handleEvent()` method in the `com.ibm.websphere.scheduler.NotificationSink` remote interface. The `handleEvent()` method is called when the notification is fired. The `Home` and `Remote` interfaces can be set as follows in the bean's deployment descriptor:

```
com.ibm.websphere.scheduler.NotificationSinkHome
com.ibm.websphere.scheduler.NotificationSink
```

The notification sink bean must exist in the same application (EAR file) that is used to create the task.

The `NotificationSink` interface defines the following method:

```
public void handleEvent(TaskNotificationInfo task) throws java.rmi.RemoteException;
```

The transactional context used by the session bean is defined by the assembler.

2. Specify the notification sink session bean to be used as the callback.

The following code illustrates how to set this option:

```
TaskInfo taskInfo = ...
Object o = new InitialContext().lookup("java:comp/env/ejb/NotificationSink");
NotificationSinkHome home =
    (NotificationSinkHome )javax.rmi.PortableRemoteObject.narrow
    (o,NotificationSinkHome.class);
taskInfo.setNotificationSink(home,TaskNotificationInfo.ALL_EVENTS);
```

3. Specify the event mask.

The event mask is specified as an integer mask. You can either use an individual mask such as `TaskNotificationInfo.CREATED` to receive specific events, `TaskNotificationInfo.ALL_EVENTS` to receive all events or a combination of specific events.. For example:

```
int eventMask = TaskNotificationInfo.CREATED+TaskNotificationInfo.PURGED;
```

## Submitting a task to a scheduler

### Before you begin

This task assumes that you have already configured a scheduler instance using the administrative console.

Once you have developed a `TaskInfo` object that contains all relevant data for a task, submit the task to a scheduler instance for creation. For example:

```
//lookup the scheduler instance to be used
Scheduler scheduler =
    (Scheduler)new InitialContext.lookup("java:comp/env/Scheduler");

TaskStatus status = scheduler.create(taskInfo);
```

**5.0.1** You must plan ahead when you create tasks with the Scheduler. When the Scheduler creates a task, the security context of the creator is stored with the scheduled task. When the task is later executed, the original context is reapplied to the thread before calling the customer `TaskInfo` instance. If you intend to secure

your application using the global security mechanism built into WebSphere Application Server, create each task with the correct credentials on the thread. Once each task has the correct credentials, you can disable and re-enable security without causing any security problems. If you do not set the security context when the Scheduler task is created and you later enable security in the target application, a security exception or error message might display.

**5.0.1** The security context is not set if any of the follow conditions are true:

- Global security is disabled
- The security context policy is disabled on the configured WorkManager for the associated Scheduler instance configuration
- A credential is not set on the thread. For example, the enterprise bean or servlet that is used to create the scheduled task is not secured.

**5.0.1** If any of the previously mentioned conditions are true when you create your task and you need to enable security on your application server or application, you must complete the following steps:

1. Cancel the task with security disabled
2. Re-create each scheduled task with security enabled, which enables Scheduler to store the desired security context with the task.

When you configure a scheduler, it is bound to a global JNDI name. Although the desired scheduler instance can be found by performing a lookup on that JNDI name, it is better to create a resource reference, which allows for more flexibility in configuring the scheduler.

**Note:** The scheduler interface is a local interface. It can only be used by server-side code; that is, J2EE applications.

Once the call to the create() method is executed, the task exists in the persistent store and is run at the time specified in the TaskInfo object. This call is transactional. If a transactional context is present on the thread when the create() method rolls back or is aborted, the task does not run.

The status object, which has been returned by the call to the create() method, contains information about the state of the task, as well as the task ID. The task ID is the unique identifier for this task, and is required if the task is to be suspended, resumed, cancelled, and so on, at a later time.

**Note:** The status object is only a snapshot of the current state of the task. Use the Scheduler.getStatus() method to receive the current state when needed.

## Managing tasks with a scheduler

When a task is created by calling the create() method on a scheduler instance, a TaskStatus object is returned to the caller. The status object contains the task ID, which is a unique identifier. The scheduler API defines several additional methods that pertain to the management of tasks, each of which accepts the task ID as a parameter. The following task management methods are defined:

### suspend()

Suspends a task. The task does not run until it has been resumed.

### resume()

Resumes a previously suspended task.

**cancel()**

Cancels a task. The task is not run.

**purge()**

Permanently deletes the task from the persistent store.

**getStatus()**

Returns the current status of the task.

For example, the following code creates and cancels a task:

```
//Create the task.
TaskInfo taskInfo = ...
TaskStatus status = scheduler.create(taskInfo);

//Get the task ID
String taskId = status.getTaskId();

//Cancel the task. Specify the purgeAlso flag so that the task does not remain
// in the persistent store
scheduler.cancel(taskId,true);
```

**Transactionality.** All methods of the scheduler API are transactional. If a global transactional context is present, it is used to perform the operation. If an unexpected exception is thrown, the transaction is marked to roll back. If an expected or declared exception is thrown, the transaction remains intact and the caller must choose to roll back or to commit the transaction. If the transaction is rolled back at some point, all scheduler operations performed within the transaction were also rolled-back.

If a local transactional context is present, it is suspended and a new global transactional context begins. Likewise, if no transactional context is active, a global transactional context begins. In both cases, if an unexpected exception is thrown, the transaction rolls back. If a declared exception is thrown, the transaction is committed.

If another thread is concurrently modifying the task in question, a `TaskPending` exception is thrown. This is because schedulers lock the database optimistically. The calling application can then retry the operation.

All methods defined by the scheduler API are described in Javadoc ([../javadoc/ee/com/ibm/websphere/scheduler/package-frame.html](http://../javadoc/ee/com/ibm/websphere/scheduler/package-frame.html)).

## Transactions and the scheduler service

### Transactions and the scheduler daemon

Scheduled `BeanTaskInfo` and `MessageTaskInfo` objects are guaranteed to execute only once. This is accomplished by grouping all of the work done in the task as a single unit of work. When each task fires, the following events occur in a single global transactional context:

1. The context of the application that created the task is applied to the thread.
2. A global transactional context is started.
3. The next fire time and start-by time are calculated using the `UserCalendar` bean or the `DefaultUserCalendar`.
4. The task database task record is updated in the database with the state of the next task or deleted if the task is complete and the task's auto-purge setting is true.
5. The `BeanTaskInfo` or `MessageTaskInfo` object is executed.

6. If the task fails and the NotificationSink bean is set, a FIRE\_FAILED notification is fired on a separate transaction.
7. If the task's NotificationSink bean is set, then the various notifications are fired as required.
8. The global transaction is committed.

Because all of a task's events are executed in a single global transactional context, you need to consider the following points in order to avoid transaction-related errors:

- Each resource participating in the task's transaction must be 2-phase XA capable.  
This includes the JDBC datasource configured for the scheduler, any JMS services used by the MessageTaskInfo objects, and any resources used within any of the UserCalendar, TaskHandler, or NotificationSink beans that have a transaction setting of "Requires".
- One resource can be single-phase, if last participant support is enabled for the application that created the transaction.  
Enable last participant support using the Application Assembly Tool. On the **WAS Enterprise** tab for your enterprise application, select the **Accept heuristic hazard** checkbox.

All unexpected exceptions are logged to the activity log and all events participating in the task's global transaction are rolled back. This includes changes to the task's database record, which force the task to be executed again when the scheduler daemon polls the database during the next poll cycle. The UserCalendar, TaskHandler, and NotificationSink beans can choose not to participate in the global transaction by setting the bean's transaction setting to "Requires new".

### Transactions and the scheduler interface

All Scheduler interface methods participate in a single global transactional context. If a global transactional context is already present on the thread when the create(), suspend(), resume(), cancel(), and purge() methods are executed, the existing global transaction is used. Otherwise, a new global transaction begins.

If the method participates in the caller's global transaction and an unexpected error occurs, the transaction is marked to roll back. If the exception is a declared exception, then the exception is rethrown to the caller, and the transaction is left alone for the caller to commit or roll back.

If the method starts its own global transaction and any exception occurs, the transaction is rolled-back, and the exception is rethrown to the caller.

## Scheduler interface

A scheduler object exists in the JNDI namespace for each scheduler configuration. A reference to a scheduler can be obtained by performing a lookup on the JNDI name; however, the lookup is valid only from the server process where the scheduler instance exists. Once a reference has been obtained, tasks can be created, suspended, cancelled, and so on, if the caller has access to the scheduler instance.

For details, see the Interface Scheduler ([../javadoc/ee/com/ibm/websphere/scheduler/Scheduler.html](http://../javadoc/ee/com/ibm/websphere/scheduler/Scheduler.html)) in the Javadoc.

### Task creation

The task is created in the persistent store using the caller's global transactional context if present. See the topic "Transactions and the scheduler service" for more details. Since this is a transactional operation, the task cannot be run or modified from another thread until the current transaction commits.

### Task modification

Tasks that have been created can be modified with the `suspend()`, `resume()`, `cancel()`, and `purge()` methods. These methods take a Task Identifier string as a parameter, which is generated by the `create()` method and can be found in the `TaskStatus` object. If a task is currently running or being modified by another thread, an operation that attempts to modify the state of the task does not block on the attempt, but a `TaskPending` exception is thrown. The operation can be reattempted at another time. Tasks can only be modified by the same application (EAR file) that was used to create the task.

### Task execution

Tasks are executed in the thread pool specified by the configuration's work manager, under the security ID of the task creator. If multiple schedulers are configured to share the same database table, the tasks found in the table can be executed on any of the schedulers, whether or not they are in the same server, node, or cell.

### Task lookup

Tasks can be located using the Name property that was assigned at creation time. This is useful when you need to modify a group of tasks and tracking individual task ID's is not convenient.

### TaskInfo interface

`TaskInfo` objects contain the information that can be used to create a task. Several implementations of this class exist, one for each type of task that can be run. Available `TaskInfo` implementations include:

#### BeanTaskInfo

Calls a stateless session bean.

#### MessageTaskInfo

Sends a JMS message to a queue or publishes a message to a topic.

For details, see the Interface `TaskInfo`

([../javadoc/ee/com/ibm/websphere/scheduler/TaskInfo.html](http://../javadoc/ee/com/ibm/websphere/scheduler/TaskInfo.html)) in the Javadoc.

After a `TaskInfo` object is created, it can be submitted to the scheduler for task creation by calling the `Scheduler.create()` method.

**Specifying time intervals.** `setStartTimeInterval()`, `setStartByInterval()`, and `setRepeatInterval()` methods all take a String parameter that represents time interval. Time intervals are calculated using *user calendars*.

### TaskHandler interface

A task handler is a user-defined stateless session bean that is called by tasks created using a `BeanTaskInfo` object. A task handler bean uses the following home and remote interfaces, which are defined in the deployment descriptor using the Application Assembly Tool or WebSphere Studio Application Developer:

```
com.ibm.websphere.scheduler.TaskHandlerHome  
com.ibm.websphere.scheduler.TaskHandler
```



The bean itself needs to implement the `process()` method defined in the remote interface. For details, see the Interface `TaskHandler` (`../javadoc/ee/com/ibm/websphere/scheduler/TaskHandler.html`) in the Javadoc.

If a task is created using a `BeanTaskInfo` object, the `process()` method on the `TaskHandler` session bean is called whenever the task runs. Because the `TaskStatus` object for the task is passed as a parameter, the handler can make use of the saved `UserContext` field, as well as determine information about the task, such as when it will fire next, the number of repeats remaining, and so on.

### **NotificationSink interface**

A notification sink is a user-defined stateless session bean that is called by tasks when their state changes throughout the bean's lifecycle. A notification sink bean uses the following home and remote interfaces, which are defined in the deployment descriptor using the Application Assembly Tool or WebSphere Studio Application Developer::

```
com.ibm.websphere.scheduler.NotificationSinkHome  
com.ibm.websphere.scheduler.NotificationSink
```

The bean itself needs to implement the `handleEvent()` method defined in the remote interface. For details, see the Interface `NotificationSink` (`../javadoc/ee/com/ibm/websphere/scheduler/NotificationSink.html`) in the Javadoc.

A `NotificationSink` provides an event notification callback on a task-by-task basis. A notification sink is set on the `TaskInfo` interface, using the `setNotificationSink()` method. If a notification sink is not specified on a task, all notifications are lost; however, the status of a task can be determined by calling the `getStatus()` method from the `Scheduler` interface. A notification callback is made for each of the following events:

- Scheduled
- Suspended
- Resumed
- Fired
- Fire Failed
- Complete
- Purged

### **UserCalendar interface**

A user calendar is a user-defined stateless session bean that is called by tasks when they need to calculate date-related values. A user calendar bean uses the following home and remote interfaces, which are defined in the deployment descriptor using the Application Assembly Tool or WebSphere Studio Application Developer:

```
com.ibm.websphere.scheduler.UserCalendarHome  
com.ibm.websphere.scheduler.UserCalendar
```

The bean itself needs to implement the `applyDelta()` and `validate()` methods defined in the remote interface. For details, see the Interface `UserCalendar` (`../javadoc/ee/com/ibm/websphere/scheduler/UserCalendar.html`) in the Javadoc.

User calendars can be used to calculate time intervals, such as the time between when a repeating task fires and the next time it fires. A user calendar takes a

java.util.Date object and applies the interval string. The resulting object is a java.util.Date object that is an incremented date.

User calendars are set by the setUserCalendar() method on the TaskInfo interface, and called by the scheduler run-time code when a delta calculation is necessary.

The following methods on the TaskInfo interface specify delta strings that use the user calendar for calculation:

- setStartTimeInterval
- setStartByInterval
- setRepeatInterval

#### **Default user calendar**

If a user calendar has not been specified using the TaskInfo.setUserCalendar() method, a default user calendar is used. The default calendar allows for simple delta specifications, such as seconds, minutes, hours, days, and months. See the JavaDoc ([../javadoc/ee/com/ibm/websphere/scheduler/UserCalendar.html](http://../javadoc/ee/com/ibm/websphere/scheduler/UserCalendar.html)) for details on the default calendar.

#### **Calendar specifiers**

A single user calendar can contain logic for multiple calendars. Which calendar is used is determined by a string that acts as the specifier. For example, a bean might be implemented to recognize the interval "day", with a specifier that determines whether to calculate "day" as a standard calendar day, or as a business day.

#### **Internationalization and timezones**

Scheduler makes use of the java.util.Date class when storing and processing dates. Internally, this class saves the time as milliseconds since the Epoch, Greenwich Mean Time. Since the Date is not converted to local time until converted to a string, scheduler respects the timezone where the date was created.

#### **Writing user calendars**

Because the user calendar is a stateless session bean, the same J2EE Programming model available to other session beans is available to the user calendar as well.

---

## **Interoperating with the Scheduler service**

### Before you begin

The WebSphere Application Server Enterprise Edition Scheduler Version 5.0 service has several interoperability problems that affect how the Scheduler behaves in a clustered environment with mixed WebSphere versions and/or mixed platforms. Specifically, there are two issues that cause interoperability problems:

1. The EJB HomeHandle format has changed between the Versions 5.0 and 5.0.1 of the product. See interim fix PQ72184 or the Interoperability section of the Version 5.0.1 Release Notes.
2. The Asynchronous Beans serialization mechanism has changed between Versions 5.0 and 5.0.1 of the product. See interim fix PQ72742 or the Asynchronous Beans section of the Version 5.0.1 Release Notes.

To avoid problems interoperating between Version 5.0 and 5.0.1 and or higher, all Version 5.0 Scheduler customers must first apply the Version 5.0 fix packs

associated with interim fixes PQ72184 and PQ72742 and remove all scheduled tasks that were created prior to applying these fix packs.

Be aware that removing all scheduled tasks involves removing data from your Scheduler databases. Follow these procedures for each Scheduler instance defined on the affected server.

#### Steps for this task

1. Recreating all Scheduler tasks.  
Use this method if you wish to programmatically recreate your scheduled tasks.
2. Deleting all Scheduler tasks.  
Use this method to remove all of your Scheduler tasks and recreate them manually.
3. Recreating all Scheduler tables.  
Use this method to drop and recreate the Scheduler tables.

## Recreating Scheduler tasks

#### Before you begin

This step requires advanced knowledge of developing J2EE Applications and the Scheduler programming interfaces.

#### Steps for this task

1. Use the administrative console to locate each Scheduler resource that was used in WebSphere Application Server Version 5.0. For each Scheduler resource, note the JNDI name.
2. Back up the Scheduler database.
3. Create a new or modify an existing EJB or servlet J2EE application to implement as a method that you intend to use as the Scheduler update program.
4. Create a method similar to the included example.

This action finds all existing Scheduler tasks, deletes them and creates new ones with the same parameters. *If Global Security is enabled in the WebSphere Application Server and the "security" context is enabled on the WorkManager referenced by the Scheduler resource, the current security context is used (as well all other J2EE contexts that are enabled on the creating EJB or servlet thread).* See the following as an example:

```
public void recreateTasks(String schedulerJNDIName)
    throws Exception
{
    InitialContext ctx = new InitialContext();
    Scheduler s = (Scheduler)ctx.lookup(schedulerJNDIName);

    Iterator tasks = null;;
    try
    {
        tasks = s.findTasksByName("%");
    }
    catch (SchedulerNotAvailableException e)
    {
        e.printStackTrace();
        throw e;
    }
}
```

```

// Iterate through each task and recreate it.
while(tasks.hasNext())
{
    TaskInfo curTask = (TaskInfo) tasks.next();

    int retries=0;
    boolean deleted=false;
    TaskStatus status=null;

    // It's best to include each cancel/create
    // in it's own transaction (not shown here).
    while(!deleted && retries < 5)
    {
        try
        {
            // Delete the task.
            s.cancel(curTask.getTaskId(), true);
            deleted = true;

            // Create a new one.
            int createRetries = 0;
            boolean created = false;
            while(!created && createRetries<5)
            {
                try
                {
                    s.create(curTask);
                    created = true;
                }
                catch (Exception e)
                {
                    ++createRetries;
                    Thread.sleep(5000);
                }
            }
            if (!created)
            {
                System.out.println("Task Not Created: " + curTask.getTaskId());
            }
        }
        catch (Exception se)
        {
            ++retries;
            Thread.sleep(5000);
        }
    }
    if (!deleted)
    {
        System.out.println("Task Not Deleted: " + curTask.getTaskId());
    }
}
}
}

```

## Deleting Scheduler tasks

### Before you begin

Use this method to remove *all* of your Scheduler tasks and recreate them manually.

### Steps for this task

1. Use the administrative console to locate each Scheduler resource that was used in WebSphere Application Server Version 5.0. For each Scheduler resource, note the following:
  - JDBC DataSource JNDI Name

- Table Prefix
2. Use the administrative console to locate each JDBC DataSource, and note the connection and location parameters.
  3. Back up the Scheduler database.
  4. Using the native database utilities, remove all of the records for each Scheduler resource for each database noted in steps 1 and 2.

The following SQL query template can be used to remove all records. Substitute TBLPFX with the actual Table Prefix.

```
DELETE FROM<TBLPFX>TASK
```

#### What to do next

Once the tasks are deleted, new tasks can be created immediately. In addition, the Scheduler service can be active while executing the DELETE SQL command in step 4. However, the operation might take longer to execute or even time out if the Scheduler has database records locked.

## Recreating Scheduler tables

#### Before you begin

Use this method to drop and recreate the Scheduler tables.

#### Steps for this task

1. Use the administrative console to locate each Scheduler resource that was used in WebSphere Application Server Version 5.0. For each Scheduler resource, note the following:
  - JDBC DataSource JNDI Name
  - Table Prefix
2. Use the administrative console to locate each JDBC DataSource, and note the connection and location parameters.
3. Back up the Scheduler database.
4. For each database and Table Prefix identified in the previous steps, execute the dropSchemaXXX.ddl and createSchemaXXX.ddl scripts located in the <install\_root>/Scheduler directory of your WebSphere Application Server, where <install\_root>, is the directory where the WebSphere Application Server is installed.

Each DDL script has instructions (located at the beginning of the respective script and in the form of a comment) on how to edit and execute the script. Refer to the Creating the database for scheduler article for more information.



---

## Chapter 21. Using shared work areas

The WorkArea service enables application developers to implicitly propagate information beyond the information passed in remote calls. Applications can create a work area, insert information into it, and make remote invocations. The work area is propagated with each remote method invocation, eliminating the need to explicitly include an appropriate argument in the definition of each method. The methods on the server side can use or ignore the information in the work area as appropriate.

Before proceeding with the steps to implement work areas, as described below, review the topic WorkArea service: Overview.

### Steps for this task

1. Developing applications that use work areas.  
Applications interact with the WorkArea service by implementing the UserWorkArea interface.
2. Managing work areas.  
The WorkArea service is managed using the administrative console.

---

## WorkArea service - Overview

One of the foundations of distributed computing is the ability to pass information, typically in the form of arguments to remote methods, from one process to another. When application-level software is written over middleware services, many of the services rely on information beyond that passed in the application's remote calls. Such services often make use of the implicit propagation of private information in addition to the arguments passed in remote requests; two typical users of such a feature are security and transaction services. Security certificates or transaction contexts are passed without the knowledge or intervention of the user or application developer. The implicit propagation of such information means that application developers do not have to manually pass the information in method invocations, which makes development less error-prone, and the services requiring the information do not have to expose it to application developers. Information such as security credentials can remain secret.

The WorkArea service gives application developers a similar facility. Applications can create a work area, insert information into it, and make remote invocations. The work area is propagated with each remote method invocation, eliminating the need to explicitly include an appropriate argument in the definition of every method. The methods on the server side can use or ignore the information in the work area as appropriate. If methods in a server receive a work area from a client and subsequently invoke other remote methods, the work area is transparently propagated with the remote requests. When the creating application is done with the work area, it terminates it.

There are two prime considerations in deciding whether to pass information explicitly as an argument or implicitly by using a work area. These considerations are:

- Pervasiveness: Is the information used in a majority of the methods in an application?
- Size: Is it reasonable to send the information even when it will not be used?

When information is sufficiently pervasive that it is easiest and most efficient to make it available everywhere, application programmers can use the WorkArea service to simplify programming and maintenance of code. The argument does not need to go onto every argument list. It is much easier to put the value into a work area and propagate it automatically. This is especially true for methods that simply pass the value on but do nothing with it. Methods that make no use of the propagated information simply ignore it.

Work areas can hold any kind of information, and they can hold an arbitrary number of individual pieces of data, each stored as a property.

## Work area property modes

The information in a work area consists of a set of properties; a property consists of a key-value-mode triple. The key-value pair represents the information contained in the property; the key is a name by which the associated value is retrieved. The mode determines whether the property can be removed or modified.

### Property modes

There are four possible mode values for properties, as shown in the following code example:

#### Code example: The PropertyModeType definition

```
public final class PropertyModeType {
    public static final PropertyModeType normal;
    public static final PropertyModeType read_only;
    public static final PropertyModeType fixed_normal;
    public static final PropertyModeType fixed_readonly;
};
```

A property's mode determines three things:

- Whether the value associated with the key can be modified
- Whether the property can be deleted
- Whether the mode associated with the key-value pair can be modified

The two read-only modes forbid changes to the information in the property; the two fixed modes forbid deletion of the property.

The WorkArea service does not provide methods specifically for the purpose of modifying the value of a key or the mode associated with a property. To change information in a property, applications simply rewrite the information in the property; this has the same effect as updating the information in the property. The mode of a property governs the changes that can be made. Modifying key-value pairs describes the restrictions each mode places on modifying the value and deleting the property. Changing modes describes the restrictions on changing the mode.

### Changing modes

The mode associated with a property can be changed only according to the restrictions of the original mode. The read-only and fixed read-only properties do not permit modification of the value or the mode. The fixed normal and fixed read-only modes do not allow the property to be deleted. This set of restrictions leads to the following permissible ways to change the mode of a property within the lifetime of a work area:



- If the current mode is normal, it can be changed to any of the other three modes: fixed normal, read-only, fixed read-only.
- If the current mode is fixed normal, it can be changed only to fixed read-only.
- If the current mode is read-only, it can be changed only by deleting the property and re-creating it with the desired mode.
- If the current mode is fixed read-only, it cannot be changed.
- If the current mode is not normal, it cannot be changed to normal. If a property is set as fixed normal and then reset as normal, the value is updated but the mode remains fixed normal. If a property is set as fixed normal and then reset as either read-only or fixed read-only, the value is updated and the mode is changed to fixed read-only.

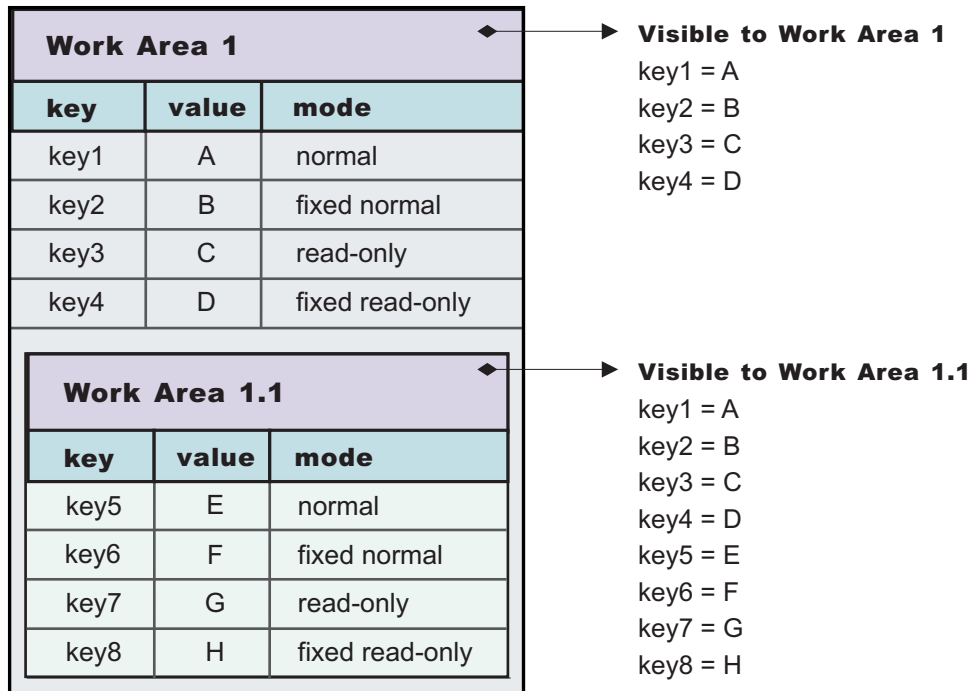
**Note:** The key, value, and mode of any property can be effectively changed by terminating (completing) the work area in which the property was created and creating a new work area. Applications can then insert new properties into the work area. This is not precisely the same as changing the value in the original work area, but some applications can use it as an equivalent mechanism.

## Nested work areas

Applications can nest work areas. When an application creates a work area, a work area context is associated with the creating thread. If the application thread creates another work area, the new work area is nested within the existing work area and becomes the current work area. Nested work areas allow applications to define and scope properties for specific tasks without having to make them available to all parts of the application. All properties defined in the original, enclosing work area are visible to the nested work area. The application can set additional properties within the nested work area that are not part of the enclosing work area.

An application working with a nested work area does not actually see the nesting of enclosing work areas. The current work area appears as a flat set of properties that includes those from enclosing work areas. In the figure below, the enclosing work area holds several properties and the nested work area holds additional properties. From the outermost work area, the properties set in the nested work area are not visible. From the nested work area, the properties in both work areas are visible.

### Defining new properties in nested work areas



Nesting can also affect the apparent settings of the properties. Properties can be deleted from or directly modified only within the work areas in which they were set, but nested work areas can also be used to temporarily override information in the property without having to modify the property. Depending on the modes associated with the properties in the enclosing work area, the modes and the values of keys in the enclosing work area can be overridden within the nested work area.

The mode associated with a property when it is created determines whether nested work areas can override the property. From the perspective of a nested work area, the property modes used in enclosing work areas can be grouped as follows:

- Modes that permit a nested work area to override the mode or the value of a key locally. The modes that permit overriding are:
  - Normal
  - Fixed normal
- Modes that do not permit a nested work area to override the mode or the value of a key locally. The modes that do not permit overriding are:
  - Read-only
  - Fixed read-only

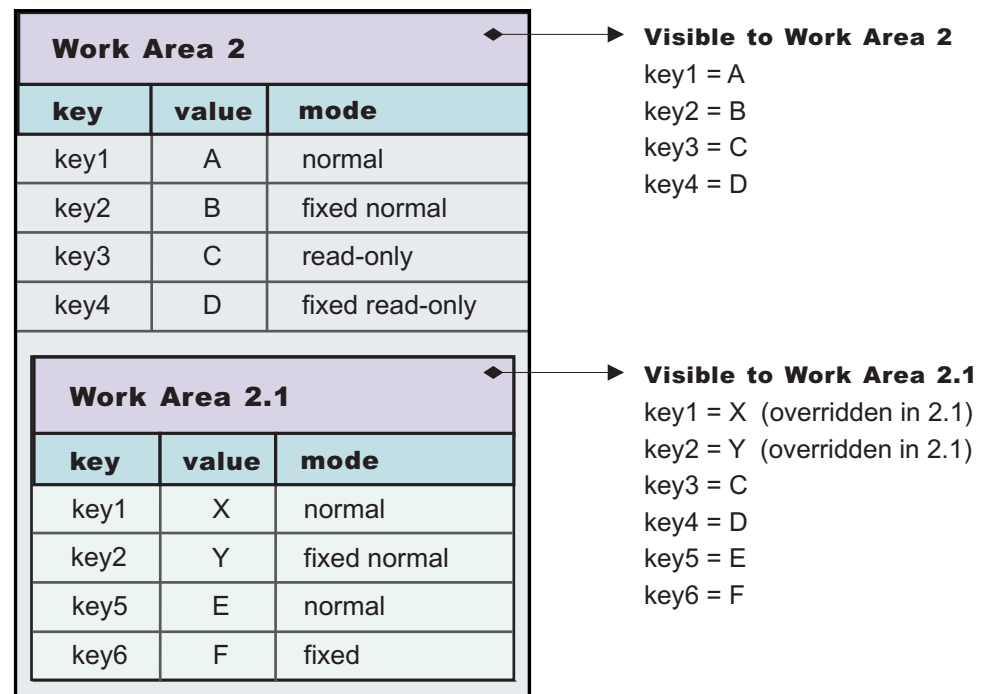
If an enclosing work area defines a property with one of the overridable modes, a nested work area can specify a new value for the key or a new mode for the property. The new value or mode becomes the value or mode seen by subsequently nested work areas. Changes to the mode are governed by the restrictions described in Changing modes. If an enclosing work area defines a property with one of the modes that cannot be overridden, no nested work area can specify a new value for the key.

A nested work area can delete properties from enclosing work areas, but the changes persist only for the duration of the nested work area. When the nested

work area is completed, any properties that were added in the nested area vanish and any properties that were deleted from the nested area are restored.

The following figure illustrates the overriding of properties from an enclosing work area. The nested work area redefines two of the properties set in the enclosing work area. The other two cannot be overridden. The nested work area also defines two new properties. From the outermost work area, the properties set or redefined in the nested work area are not visible. From the nested work area, the properties in both work areas are visible, but the values seen for the redefined properties are those set in the nested work area.

### Redefining existing properties in nested work areas



## Distributed work areas

If a remote invocation is issued from a thread associated with a work area, a copy of the work area is automatically propagated to the target object, which can use or ignore the information in the work area as necessary. If the calling application has a nested work area associated with it, a copy of the nested work area and all its ancestors is propagated to the target. The target application can locally modify the information, as allowed by the property modes, by creating additional nested work areas; this information will be propagated to any remote objects it invokes. However, no changes made to a nested work area on a target object are propagated back to the calling object. The caller's work area is unaffected by changes made in the remote method.

## WorkArea service: Special considerations

Developers who use work areas should consider the following issues that could potentially cause problems: interoperability between the EJB and CORBA programming models; and the use of work areas with Java's Abstract Windowing Toolkit.

## EJB and CORBA interoperability

Although the WorkArea service can be used across the EJB and CORBA programming models, many composed data types cannot be successfully used across those boundaries. For example, if a SimpleSampleCompany instance is passed from the WebSphere environment into a CORBA environment, the CORBA application can retrieve the SimpleSampleCompany object encapsulated within a CORBA Any object from the work area, but it cannot extract the value from it. Likewise, an IDL-defined struct defined within a CORBA application and set into a work area will not be readable by an application using the UserWorkArea class. Applications can avoid this incompatibility by directly setting only primitive types, like integers and strings, as values in work areas, or by implementing complex values with structures designed to be compatible, like CORBA valuetypes. Also, CORBA Anys that contains either the tk\_null or tk\_void typecode can be set into the work area by using the CORBA interface, but the work-area specification cannot allow the J2EE implementation to return null on a lookup that retrieves these CORBA-set properties without incorrectly implying that there is no value set for the corresponding key. If a J2EE application tries to retrieve CORBA-set properties that are non-serializable, or contain CORBA nulls or void references, the `com.ibm.websphere.workarea.IncompatibleValue` exception is raised.

### Using work areas with Java's Abstract Windowing Toolkit (AWT)

Work areas must be used cautiously in applications that use Java's Abstract Windowing Toolkit (AWT). The AWT implementation is multithreaded, and work areas begun on one thread are not available on another. For example, if a program begins a work area in response to an AWT event, such as pressing a button, the work area might not be available to any other part of the application after the execution of the event completes.

## WorkArea service performance considerations

The WorkArea service is designed to address complex data passing patterns that can quickly grow beyond convenient maintenance. A work area is a note pad that is accessible to any client capable of looking up Java Naming Directory Interface (JNDI). Once a work area is established, data can be placed there for future use in any subsequent method calls to both remote and local resources.

You can utilize a work area when a large number of methods require common information or if information is only needed by a method that is significantly further down the call graph. The former avoids the need for complex parameter passing models where the number of arguments passed becomes excessive and hard to maintain. You can improve application function by placing the information in a work area and subsequently accessing it independently in each method, eliminating the need to pass these parameters from method to method. The latter case also avoids unnecessary parameter passing and helps to improve performance by reducing the cost of marshalling and de-marshalling these parameters over the Object Request Broker (ORB) when they are only needed sparsely throughout the call graph.

When attempting to maximize performance by using a work area, cache the UserWorkArea retrieved from JNDI wherever it is accessed. You can reduce the amount of time spent looking up information in JNDI by retrieving it once and keeping a reference for the future. JNDI look up takes time and can be costly.

Performance is degraded if you use a work area as a direct replacement to passing a single parameter over a single method call. The reason is that you incur more overhead than just passing that parameter between method calls. Although the degradation is usually within acceptable tolerances and scales similarly to passing parameters with regard to object size, consider this potential problem before utilizing the service. As with most functional services, intelligent usage of work areas yields the best results.

The WorkArea service is a tool to simplify the job of passing information from resource to resource, and in some cases can improve performance by reducing the overhead associated with parameter passing when the information is only sparsely accessed within the call graph. Caching the instance retrieved from JNDI is important to effectively maximize performance during runtime.

---

## Developing applications that use work areas

Applications interact with the WorkArea service by implementing the UserWorkArea interface. This interface defines all of the methods used to create, manipulate, and complete work areas:

### Steps for this task

1. Accessing the WorkArea service.
2. Beginning a work area.
3. Setting properties in a work area.
4. Using a work area to manage local work.
5. Completing a work area.

### Usage scenario

An example application, the WorkArea SimpleSample application, is used throughout this documentation to illustrate these tasks

## UserWorkArea interface

Applications interact with the WorkArea service by implementing the UserWorkArea interface. This interface, shown below, defines all of the methods used to create, manipulate, and terminate work areas:

```
package com.ibm.websphere.workarea;

public interface UserWorkArea {
    void begin(String name);
    void complete() throws NoWorkArea, NotOriginator;

    String getName();
    String[] retrieveAllKeys();
    void set(String key, java.io.Serializable value)
        throws NoWorkArea, NotOriginator, PropertyReadOnly;
    void set(String key, java.io.Serializable value, PropertyModeType mode)
        throws NoWorkArea, NotOriginator, PropertyReadOnly;
    java.io.Serializable get(String key);
    PropertyModeType getMode(String key);
    void remove(String key)
        throws NoWorkArea, NotOriginator, PropertyFixed;
}
```

**Note:** EJB applications can use the UserWorkArea interface only within the implementation of methods in the remote interface; likewise, servlets can use the

interface only within the service method of the HttpServlet class. Use of work areas within any lifecycle method of a servlet or enterprise bean is considered a deviation from the work area programming model and is not supported.

### Exceptions

The WorkArea service defines the following exceptions for use with the UserWorkArea interface:

#### NoWorkArea

Thrown when a request requires an associated work area but none is present.

#### NotOriginator

Raised when a request attempts to manipulate the contents of an imported work area.

#### PropertyReadOnly

Raised when a request attempts to modify a read-only or fixed read-only property.

#### PropertyFixed

Raised by the remove method when the designated property has one of the fixed modes.

## Example: WorkArea SimpleSample application

In this example, the client creates a work area and inserts two properties into the work area: a site identifier and a priority. The site-identifier is set as a read-only property; the client does not allow recipients of the work area to override the site identifier. This property consists of the key company and a static instance of a SimpleSampleCompany object. The priority property consists of the key priority and a static instance of a SimpleSamplePriority object. The object types are defined as shown in the following code example

```
public static final class SimpleSampleCompany {
    public static final SimpleSampleCompany Main;
    public static final SimpleSampleCompany NewYork_Sales;
    public static final SimpleSampleCompany NewYork_Development;
    public static final SimpleSampleCompany London_Sales;
    public static final SimpleSampleCompany London_Development;
}

public static final class SimpleSamplePriority {
    public static final SimpleSamplePriority Platinum;
    public static final SimpleSamplePriority Gold;
    public static final SimpleSamplePriority Silver;
    public static final SimpleSamplePriority Bronze;
    public static final SimpleSamplePriority Tin;
}
```

The client then makes an invocation on a remote object. The work area is automatically propagated; none of the methods on the remote object take a work area argument. On the remote side, the request is first handled by the SimpleSampleBean; the bean first reads the site identifier and priority properties from the work area. The bean then intentionally attempts, and fails, both to write directly into the imported work area and to override the read-only site-identifier property.

The SimpleSampleBean successfully begins a nested work area, in which it overrides the client's priority, then calls another bean, the SimpleSampleBackendBean. The SimpleSampleBackendBean reads the properties

from the work area, which contains the site identifier set in the client and priority set in the SimpleSampleBean. Finally, the SimpleSampleBean completes its nested work area, writes out a message based on the site-identifier property, and returns.

The implementation of this application is discussed in the topic, Developing applications that use work areas.

## Accessing the WorkArea service

The WorkArea service provides a JNDI binding to an implementation of the UserWorkArea interface under the name `java:comp/websphere/UserWorkArea`. Applications that need to access the service can perform a lookup on that JNDI name, as shown in the following code example:

### Usage scenario

```
import com.ibm.websphere.workarea.*;
import javax.naming.*;

public class SimpleSampleServlet {
    ...

    InitialContext jndi = null;
    UserWorkArea userWorkArea = null;
    try {
        jndi = new InitialContext();
        userWorkArea = (UserWorkArea)jndi.lookup(
            "java:comp/websphere/UserWorkArea");
    }
    catch (NamingException e) { ... }
}
```

### What to do next

The next step is to use the `begin()` method to create a new work area and associate it with the calling thread, as described in the topic Beginning a new work area.

## Beginning a new work area

### Before you begin

Be sure that your client has a reference to the UserWorkArea interface, as described in the topic Accessing the WorkArea Service

Use the `begin()` method to create a new work area and associate it with the calling thread. The `begin()` method takes a string as an argument; the string is used to name the work area. The argument must not be null, which causes the `java.lang.NullPointerException` to be raised. In the following code example, the application begins a new work area with the name `SimpleSampleServlet`:

```
public class SimpleSampleServlet {
    ...
    try {
        ...
        userWorkArea = (UserWorkArea)jndi.lookup(
            "java:comp/websphere/UserWorkArea");
    }
    ...

    userWorkArea.begin("SimpleSampleServlet");
    ...
}
```

The `begin()` method is also used to create nested work areas; if a work area is associated with a thread when the `begin()` method is called, the method creates a new work area nested within the existing work area.

The `WorkArea` service makes no use of the names associated with work areas; You can name work areas in any way that you choose. Names are not required to be unique, but the usefulness of the names for debugging is enhanced if the names are distinct and meaningful within the application. Applications can use the `getName()` method to return the name associated with a work area by the `begin()` method.

### What to do next

Using a work area

## Setting properties in a work area

An application with a current work area can insert properties into the work area and retrieve the properties from the work area. The `UserWorkArea` interface provides two `set()` methods for setting properties and a `get()` method for retrieving properties. The two-argument `set()` method inserts the property with the property mode of `normal`. The three-argument `set()` method takes a property mode as the third argument. (See "Setting property modes", later in this topic.)

Both `set()` methods take the key and the value as arguments. The key is a `String`; the value is an object of the type `java.io.Serializable`. None of the arguments can be null, which causes the `java.lang.NullPointerException` exception to be raised.

The `WorkArea` `SimpleSample` application uses objects of two classes, the `SimpleSampleCompany` class and the `SimpleSampleProperty` class, as values for properties. The `SimpleSampleCompany` class is used for the site identifier, and the `SimpleSamplePriority` class is used for the priority. These classes are shown in following code example:

```
public class SimpleSampleServlet {
    ...
    userWorkArea.begin("SimpleSampleServlet");

    try {
        // Set the site-identifier (default is Main).
        userWorkArea.set("company",
            SimpleSampleCompany.Main, PropertyModeType.read_only);

        // Set the priority.
        userWorkArea.set("priority",
            SimpleSamplePriority.Silver);
    }

    catch (PropertyReadOnly e) {
        // The company was previously set with the read-only or
        // fixed read-only mode.
        ...
    }

    catch (NotOriginator e) {
        // The work area originated in another process,
        // so it can't be modified here.
        ...
    }

    catch (NoWorkArea e) {
        // There is no work area begun on this thread.
    }
}
```



```

    }
    ...
    // Do application work.
    ...
}

```

The `get()` method takes the key as an argument and returns a Java Serializable object as the value associated with the key. For example, to retrieve the value of the company key from the work area, the code example above uses the `get()` method on the work area to retrieve the value.

**Setting property modes.** The two-argument `set()` method on the `UserWorkArea` interface takes a key and a value as arguments and inserts the property with the default property mode of normal. To set a property with a different mode, applications must use the three-argument `set()` method, which takes a property mode as the third argument. The values used to request the property modes are as follows:

- **Normal:** `PropertyModeType.normal`
- **Fixed normal:** `PropertyModeType.fixed_normal`
- **Read-only:** `PropertyModeType.read_only`
- **Fixed read-only:** `PropertyModeType.fixed_readonly`

## Using a work area to manage local work

### Before you begin

Be sure that your client has a reference to the `UserWorkArea` interface, as described in the topic [Accessing the WorkArea Service](#)

In a business application that uses work areas, server objects typically retrieve the work area properties and use them to guide local work.

### Steps for this task

1. Retrieving the name of the active work area  
This step determines whether the calling thread is associated with a work area.
2. Overriding work area properties.  
Server objects can override client work area properties by creating their own, nested work area.
3. Retrieving properties from a work area
4. Retrieving a list of all keys in a work area
5. Querying the mode of a work area property
6. Deleting a work area property
7. Completing a work area

### Usage scenario

The server side of the `WorkArea SimpleSample` application accepts remote invocations from clients. With each remote call, the server also gets a work area from the client if the client has created one. The work area is propagated transparently. None of the remote methods includes the work area on its argument list.

In the example application, the server objects use the work area interface for demonstration purposes only. For example, the SimpleSampleBean intentionally attempts to write directly to an imported work area, which triggers the NotOriginator exception. Likewise, the bean intentionally attempts to mask the read only SimpleSampleCompany, which triggers the PropertyReadOnly exception. The SimpleSampleBean also nests a work area and successfully overrides the priority property before invoking the SimpleSampleBackendBean. A true business application would extract the work area properties and use them to guide the local work. The SimpleSampleBean mimics this by writing a message that function is denied when a request emanates from a sales environment.

### Retrieving the name of the active work area

Applications use the getName() method on the UserWorkArea interface to retrieve the name of the current work area. This is the recommended method for determining whether the thread is associated with a work area; if the thread is not associated with a work area, the getName() method returns null. In the following code example, the name of the work area corresponds to the name of the class in which the work area was begun.

#### Usage scenario

```
public class SimpleSampleBeanImpl implements SessionBean {
    ...

    public String [] test() {
        // Get the work-area reference from JNDI.
        ...

        // Retrieve the name of the work area. In this example,
        // the name is used to identify the class in which the
        // work area was begun.
        String invoker = userWorkArea.getName();
        ...
    }
}
```

### Overriding work area properties

Work areas are inherently associated with the process that creates them. In the sample application, the client begins a work area and sets into it the site-identifier and priority properties. This work area is propagated to the server when the client makes a remote invocation.

Applications nest work areas in order to temporarily override properties imported from a client process. The nesting mechanism is automatic; invoking begin on the UserWorkArea interface from within the scope of an existing work area creates a nested work area that inherits the properties from the enclosing work area. Properties set into the nested work area are strictly associated with the process in which the work area was begun; the nested work area must be completed within the process that created them. If a work area is not completed by the creating process, the work-area facility terminates the work area when the process exits. After a nested work area is completed, the original view of the enclosing work area is restored. However, the view of the complete set of work areas associated with a thread cannot be decomposed by downstream processes.

Applications set properties into a work area using property modes in ensure that a particular property is fixed (not removable) or read-only (not overrideable) within the scope of the given work area.

#### Usage scenario

In the following code example, the server-side sample bean attempts to write directly to the imported work area; this action is not permitted, and the `NotOriginator` exception is thrown. The sample bean must begin its own work area in order to override any imported properties, as shown in the second code example.

```
public class SimpleSampleBeanImpl implements SessionBean {

    public String [] test() {
        ...
        String invoker = userWorkArea.getName();

        try {
            userWorkArea.set("key", "value");
        }
        catch (NotOriginator e) {
        }
        ...
    }
}
```

The following code example demonstrates beginning a nested work area, using the name of the creating class to identify the nested work area.

```
public class SimpleSampleBeanImpl implements SessionBean {

    public String [] test() {
        ...
        String invoker = userWorkArea.getName();
        try {
            userWorkArea.set("key", "value");
        }
        catch (NotOriginator e) {
        }

        // Begin a nested work area. By using the name of the creating
        // class as the name of the work area, we can avoid having
        // to explicitly set the name of the creating class in
        // the work area.
        userWorkArea.begin("SimpleSampleBean");

        ...
    }
}
```

In the example application, the client sets the site-identifier property as read-only; that guarantees that the request will always be associated with the client's company identity. A server cannot override that value in a nested work area. In the following code example, the `SimpleSampleBean` attempts to change the value of the site-identifier property in the nested work area it created.

```
public class SimpleSampleBeanImpl implements SessionBean {

    public String [] test() {
        ...

        String invoker = userWorkArea.getName();
        try {
            userWorkArea.set("key", "value");
        }
        catch (NotOriginator e) {
        }

        // Begin a nested work area.
        userWorkArea.begin("SimpleSampleBean");
    }
}
```

```

    try {
        userWorkArea.set("company",
                        SimpleSampleCompany.London_Development);
    }
    catch (NotOriginator e) {
    }
    ...
}

```

### Retrieving work area properties

Properties can be retrieved from a work area by using the `get()` method. This method is intentionally lightweight; there are no declared exceptions to handle. If there is no active work area, or if there is no such property set in the current work area, the `get()` method returns null.

**Note:** The `get()` method can raise a `NotSerializableError` in the relatively rare scenario in which CORBA clients set composed data types and invoke enterprise-bean interfaces.

#### Usage scenario

The following example shows the retrieval of the site-identifier and priority properties by the `SimpleSampleBean`. Recall that one property was set into an outer work area by the client, and the other property was set into the nested work area by the server-side bean; the nesting is transparent to the retrieval of the properties.

```

public class SimpleSampleBeanImpl implements SessionBean {

    public String [] test() {
        ...

        // Begin a nested work area.
        userWorkArea.begin("SimpleSampleBean");
        try {
            userWorkArea.set("company",
                            SimpleSampleCompany.London_Development);
        }
        catch (NotOriginator e) {
        }

        SimpleSampleCompany company =
            (SimpleSampleCompany) userWorkArea.get("company");
        SimpleSamplePriority priority =
            (SimpleSamplePriority) userWorkArea.get("priority");
        ...
    }
}

```

### Retrieving a list of all keys in a work area

The `UserWorkArea` interface provides the `retrieveAllKeys()` method for retrieving a list of all the keys visible from a work area. This method takes no arguments and returns an array of strings. This method returns null if there is no work area associated with the thread. If there is an associated work area containing no properties, the method returns an array of size 0.

### Querying the mode of a work area property

The `UserWorkArea` interface provides the `getMode()` method for determining the mode of a specific property. This method takes the property's key as an argument and returns the mode as a `PropertyModeType` object. (See Setting property modes for more information on names of mode types.) If the specified key does not exist

in the work area, the method returns `PropertyModeType.normal`, indicating that the property can be set and removed without error.

### Deleting a work area property

The `UserWorkArea` interface provides the `remove()` method for deleting a property from the current scope of a work area. If the property was initially set in the current scope, removing it deletes the property. If the property was initially set in an enclosing work area, removing it deletes the property until the current scope is completed. When the current work area is completed, the deleted property is restored.

The `remove()` method takes the property's key as an argument. Only properties with the modes `normal` and `read-only` can be removed. Attempting to remove a fixed property causes the `PropertyFixed` exception to be thrown. Attempting to remove properties in work areas that originated in other processes causes the `NotOriginator` exception to be thrown.

## Completing a work area

After an application has finished using a work area, it must complete the work area by calling the `complete()` method on the `UserWorkArea` interface. This terminates the association with the calling thread and destroys the work area. If the `complete` method is called on a nested work area, the nested work area is terminated and the parent work area becomes the current work area. If there is no work area associated with the calling thread, a `NoWorkArea` exception is thrown.

Every work area must be completed, and work areas can be completed only by the originating process. For example, if a server attempts to call the `complete()` method on a work area that originated in a client, a `NotOriginator` exception is thrown. Work areas created in a server process are never propagated back to an invoking client process.

**Note:** The `WorkArea` service claims full local-remote transparency. Even if two beans happen to be deployed in the same server, and therefore the same JVM and process, a work area begun on an invocation from another is completed and the bean in which the request originated is always in the same state after any remote call.

### Usage scenario

The following code example shows the completion of the work area created in the client application.

```
public class SimpleSampleServlet {
    ...
    userWorkArea.begin("SimpleSampleServlet");
    userWorkArea.set("company",
        SimpleSampleCompany.Main, PropertyModeType.read_only);
    userWorkArea.set("priority", SimpleSamplePriority.Silver);
    ...

    // Do application work.
    ...

    // Terminate the work area.
    try {
        userWorkArea.complete();
    }

    catch (NoWorkArea e) {
```

```

        // There is no work area associated with this thread.
        ...
    }

    catch (NotOriginator e) {
        // The work area was imported into this process.
        ...
    }
    ...
}

```

The following code example shows the sample application completing the nested work area it created earlier in the remote invocation.

```

public class SimpleSampleBeanImpl implements SessionBean {

    public String [] test() {
        ...

        // Begin a nested work area.
        userWorkArea.begin("SimpleSampleBean");
        try {
            userWorkArea.set("company",
                SimpleSampleCompany.London_Development);
        }
        catch (NotOriginator e) {
        }

        SimpleSampleCompany company =
            (SimpleSampleCompany) userWorkArea.get("company");
        SimpleSamplePriority priority =
            (SimpleSamplePriority) userWorkArea.get("priority");

        // Complete all nested work areas before returning.
        try {
            userWorkArea.complete();
        }
        catch (NoWorkArea e) {
        }
        catch (NotOriginator e) {
        }
    }
}

```

---

## Managing the work area service

The WorkArea service is managed using the administrative console. There are two administrative tasks associated with work areas:

- Enabling the WorkArea service. The WorkArea Service is enabled by default on both clients and servers.
- Managing the size of work areas. Applications can set maximum sizes on each work area to be sent and to be accepted.

### Enabling the WorkArea service

For an application to take advantage of work areas, the WorkArea service must be enabled for both clients and servers. In both cases, the service is enabled by default.

#### Steps for this task

1. Enable (or disable) the use of work areas on a server:

- a. Start the administrative console.
- b. Select **Servers** > *server\_name* > **WorkArea Service**.
- c. Select or clear the **Startup** checkbox.

This specifies whether or not the server should automatically start the WorkArea service when the server starts.

- 2. Enable (or disable) the use of work areas on a client:

Set the `com.ibm.websphere.workarea.enabled` property to `TRUE` or `FALSE` before starting the client. For example, edit the `launchClient` script in the `$WAS_HOME/bin` directory and add the following to the Java invocation:

```
-Dcom.ibm.websphere.workarea.enabled=false
```

### WorkArea service settings

Use this page to manage the work area service.

The work area service manages the scope and implicit propagation of application context.

To view this administrative console page, click **Servers** > **Application Servers** > *server\_name* > **Work Area Service** .

**Startup:** Specifies whether the server will attempt to start the work area service.

|         |   |
|---------|---|
| Default | Selected  |
| Range   | <p><b>Selected</b></p> <p>When the application server starts, it attempts to start the work area service automatically.</p> <p><b>Cleared</b></p> <p>The server does not try to start the work area service. If work areas are to be used on this application server, the system administrator must start the service manually or select this property then restart the server.</p> |

**Maximum Send Size:** Specifies the maximum size of data that can be sent within a single work area.

|           |                    |
|-----------|--------------------|
| Data type | Integer            |
| Units     | Bytes              |
| Default   | 32767              |
| Range     | -1 to no limit     |
|           | -1      Default.   |
|           | 0        No limit. |

**Maximum Receive Size:** Specifies the maximum size of data that can be received by a single work area.

|           |         |
|-----------|---------|
| Data type | Integer |
| Units     | Bytes   |
| Default   | 32767   |

|       |                |
|-------|----------------|
| Range | -1 to no limit |
|       | -1 Default.    |
|       | 0 No limit.    |

## Managing the size of work areas

Applications can set maximum sizes on each work area to be sent or received. By default, the maximum size of a work area that is sent by a client and received, then possibly re-sent, by a server is 32,768 bytes. You can change this size as described in this topic.

### Steps for this task

1. Change the size of the work area that can be sent or received by a server:
  - a. Start the administrative console.
  - b. Select **Servers** > *server\_name* > **WorkArea Service**.
  - c. Enter a new value in the **maxSendSize** field to modify the size of the work area that this server can send, or enter a new value in the **maxReceiveSize** field to modify the size of the work area that this server can accept.

2. Change the size of the work area that can be sent by a client:

Set the `com.ibm.websphere.workarea.maxSendSize` property to the desired number of bytes before starting the client. This can be done in several ways. For example, to set the maximum size to 10,000 bytes, edit the `launchClient` script in the `$WAS_HOME/bin` directory and add the following to the Java invocation:

```
-Dcom.ibm.websphere.workarea.maxSendSize=10000
```

### Results

The maximum size that can be specified is determined by the maximum value expressible in the Java Integer data type, 2,147,483,647. The smallest maximum size that can be specified is 1. Using a maximum size of 1 byte effectively means that no requests associated with the work area can leave the system or enter another system. A value of 0 means that no limit is imposed. A value of -1 means that the default value is to be honored. The default value is also used if an invalid value or a malformed property is specified.



---

## Chapter 22. Using the transaction service

These topics provide information about using transactions with WebSphere applications

WebSphere applications can use transactions to coordinate multiple updates to resources as atomic units (as indivisible units of work) such that all or none of the updates are made permanent.

In WebSphere Application Server, transactions are handled by three main components:

- A transaction manager that supports the enlistment of recoverable XAResources and ensures that each such resource is driven to a consistent outcome either at the end of a transaction or after a failure and restart of the application server.
- A container in which the J2EE application runs. The container manages the enlistment of XAResources on behalf of the application when the application performs updates to transactional resource managers (for example, databases). Optionally, the container can control the demarcation of transactions for enterprise beans configured for container-managed transactions.
- An application programming interface (UserTransaction) that is available to bean-managed enterprise beans and servlets. This allows such application components to control the demarcation of their own transactions.

For more information about using transactions with WebSphere applications, see the following topics:

- "Transaction support in WebSphere Application Server"
- "Developing components to use transactions"
- "Configuring transaction properties for an application server"
- "Using local transactions"
- "Setting transactional attributes in the deployment descriptor"
- "Using bean-managed transactions"
- "Managing active transactions"
- "Troubleshooting transactions"
- "Transaction service exceptions"
- "UserTransaction interface - methods available"
- Coordinating access to 1-PC and 2-PC-capable resources within the same transaction
- Implementing WebSphere enterprise applications that use ActivitySessions

---

### Transaction support in WebSphere Application Server

A transaction is unit of activity within which multiple updates to resources can be made atomic (as an indivisible unit of work) such that all or none of the updates are made permanent. For example, multiple SQL statements to a relational database are committed atomically by the database during the processing of an SQL COMMIT statement. In this case, the transaction is contained entirely within the database manager and can be thought of as a *resource manager local transaction (RMLT)*. In some contexts, a transaction is referred to as a logical unit of work (LUW).

The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself).
- Entity beans use container-managed transactions.
- Web components (servlets) use bean-managed transactions.

WebSphere Application Server is a transaction manager that supports the coordination of resource managers through their XAResource interface and participates in distributed global transactions with other OTS 1.2 compliant transaction managers (for example J2EE 1.3 application servers). WebSphere applications can also be configured to interact with databases, JMS queues, and JCA connectors through their *local transaction* support when distributed transaction coordination is not required.

Resource managers that offer transaction support can be categorized into those that support two-phase coordination (by offering an XAResource interface) and those that support only one-phase coordination (for example through a LocalTransaction interface). The WebSphere Application Server transaction support provides coordination, within a transaction, for any number of two-phase capable resource managers. It also enables a single one-phase capable resource manager to be used within a transaction in the absence of any other resource managers, although a WebSphere transaction is not necessary in this case.

With the Last Participant Support of WebSphere Application Server Enterprise, you can coordinate the use of a single one-phase commit (1PC) capable resource with any number of two-phase commit (2PC) capable resources in the same global transaction. At transaction commit, the two-phase commit resources are prepared first using the two-phase commit protocol, and if this is successful the one-phase commit-resource is then called to `commit(one_phase)`. The two-phase commit resources are then committed or rolled back depending on the response of the one-phase commit resource.

The ActivitySession service of WebSphere Application Server Enterprise provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. It is a distributed context that can be used to coordinate multiple one-phase resource managers. The WebSphere EJB container and deployment tooling support ActivitySessions as an extension to the J2EE programming model. EJBs can be deployed with lifecycles that are influenced by ActivitySession context, as an alternative to transaction context. An application can then interact with a resource manager through its LocalTransaction interface for the period of a client-scoped ActivitySession rather than just the duration of an EJB method.

## Resource manager local transaction (RMLT)

A resource manager local transaction (RMLT) is a resource manager's view of a local transaction; that is, it represents a unit of recovery on a single connection that is managed by the resource manager.

Resource managers include:

- Enterprise Information Systems that are accessed through a resource adapter, as described in the J2EE Connector Architecture 1.0 (<http://java.sun.com/j2ee/connector/index.html>).
- Relational databases that are accessed through a JDBC datasource.

- JMS queue and topic destinations.

Resource managers offer specific interfaces to enable control of their RMLTs. J2EE connector resource adapters that include support for local transactions provide a LocalTransaction interface to enable applications to request that the resource adapter commit or rollback RMLTs. JDBC datasources provide a Connection interface for the same purpose.

The boundary at which all RMLTs must be complete is defined in WebSphere Application Server by a "Local transaction containment (LTC)".

## Global transactions

If an application uses two or more resources, then an external transaction manager is needed to coordinate the updates to both resource managers in a *global transaction*.

Global transaction support is available to web and enterprise bean J2EE components. Enterprise bean components can be subdivided into beans that exploit container-managed transactions (CMT) or bean-managed transactions (BMT).

BMT enterprise beans and web components can use the Java Transaction API (JTA) UserTransaction interface to define the demarcation of a global transaction. The UserTransaction interface is obtained by a JNDI lookup of `java:comp/UserTransaction`. The UserTransaction is not available to the following components:

- CMT enterprise beans. Any attempt by such beans to obtain the interface results in an exception in accordance with the EJB specification.
- Client applications running outside the Web and EJB containers.

Ensure that programs that perform a JNDI lookup of the UserTransaction interface, use an InitialContext that resolves to a local implementation of the interface. Also ensure that such programs use a JNDI location appropriate for the EJB version.

Before the EJB 1.1 specification, the JNDI location of the UserTransaction interface was not specified. Each EJB container implementor defined it in an implementation-specific manner. Earlier versions of WebSphere Application Server, up to and including Version 3.5.x (without EJB 1.1), bind the UserTransaction interface to a JNDI location of `jta/usertransaction`. WebSphere Application Server Version 4, and later releases, bind the UserTransaction interface at the location defined by EJB 1.1, which is `java:comp/UserTransaction`. WebSphere Application Server, Version 5 no longer provides the `jta/usertransaction` binding within Web and EJB containers to applications at a J2EE level of 1.3 or later. For example, EJB 2.0 applications can use only the `java:comp/UserTransaction` location.

A web component or enterprise bean (CMT or BMT) can get the ExtendedJTATransaction interface through a lookup of `java:comp/websphere/ExtendedJTATransaction`. This interface provides access to the transaction identity and a mechanism to receive notification of transaction completion.

## Local transaction containment (LTC)

A *local transaction containment (LTC)* is used to define the application server behavior in an unspecified transaction context.

(Unspecified transaction context is defined in the Enterprise JavaBeans 2.0 Specification (<http://java.sun.com/products/ejb/2.0.html>.)

A LTC is a bounded unit-of-work scope within which zero, one, or more resource manager local transactions (RMLTs) can be accessed. The LTC defines the boundary at which all RMLTs must be complete; any incomplete RMLTs are resolved, according to policy, by the container. An LTC is local to a bean instance; it is not shared across beans even if those beans are managed by the same container. LTCs are started by the container before dispatching a method on a J2EE component (such as an enterprise bean or servlet) whenever the dispatch occurs in the absence of a global transaction context. LTCs are completed by the container depending on the application-configured LTC boundary; for example at the end of the method dispatch. There is no programmatic interface to the LTC support; rather LTCs are managed exclusively by the container and configured by the application deployer through transaction attributes in the application deployment descriptor.

A local transaction containment cannot exist concurrently with a global transaction. If application component dispatch occurs in the absence of a global transaction, the container always establishes an LTC. The only exceptions to this are as follows:

- Where application component dispatch occurs without container interposition; for example, for a stateless session bean create.
- J2EE 1.2 web components.
- J2EE 1.2 BMT enterprise beans.

A local transaction containment can be scoped to an `ActivitySession` context that lives longer than the enterprise bean method in which it is started, as described in “`ActivitySession` and transaction contexts” on page 1023 `ActivitySessions` and `transaction contexts`.

## Using local transactions

Local transaction containment (LTC) support, and its configuration through local transaction extended deployment descriptors, gives IBM WebSphere Application Server application programmers a number of advantages. This topic describes those advantages and how they relate to the settings of the local transaction extended deployment descriptors. This topic also describes points to consider to help you best configure transaction support for some example scenarios that use local transactions.

### **Develop an enterprise bean or servlet that accesses one or more databases that are independent and require no coordination.**

If an enterprise bean does not need to use global transactions, it is often more efficient to deploy the bean with the Container Transaction deployment descriptor **Transaction** attribute set to `Not supported` instead of `Required`.

With the extended local transaction support of IBM WebSphere Application Server, applications can perform the same business logic in an unspecified transaction context as they can under a global transaction. An enterprise bean, for example, runs under an unspecified transaction context if it is deployed with a **Transaction** attribute of `Not supported` or `Never`.

The extended local transaction support provides a container-managed, implicit local transaction boundary within which application updates can be committed and their connections cleaned up by the container. Applications can then be designed with a greater degree of independence from deployment concerns. This makes using a **Transaction** attribute of

Supports much simpler, for example, when the business logic may be called either with or without a global transaction context.

An application can follow a get-use-close pattern of connection usage regardless of whether or not the application runs under a transaction. The application can depend on the close behaving in the same way and not causing a rollback to occur on the connection if there is no global transaction.

There are many scenarios where ACID coordination of multiple resource managers is not needed. In such scenarios running business logic under a **Transaction** policy of Not supported performs better than if it had been run under a Required policy. This benefit is exploited through the **Local Transactions - Resolution-control** extended deployment setting of ContainerAtBoundary. With this setting, application interactions with resource providers (such as databases) are managed within implicit RMLTs that are both started and ended by the container. The RMLTs are committed by the container at the configured **Local Transactions - Boundary**; for example at the end of a method. If the application returns control to the container by an exception, the container rolls back any RMLTs that it has started.

This usage applies to both servlets and enterprise beans.

#### **Use local transactions in a managed environment that guarantees clean-up.**

Applications that want to control RMLTs, by starting and ending them explicitly, can use the default **Local Transactions - Resolution-control** extended deployment setting of Application. In this case, the container ensures connection cleanup at the boundary of the local transaction context.

J2EE specifications that describe application use of local transactions do so in the manner provided by the default setting of **Local Transactions - Resolution-control=Application** and **Local Transactions - Unresolved-action=Rollback**. By configuring the **Local Transactions - Unresolved-action** extended deployment setting to Commit, then any RMLTs started by the application but not completed when the local transaction containment ends (for example, when the method ends) are committed by the container. This usage applies to both servlets and enterprise beans.

#### **Extend the duration of a local transaction beyond the duration of an EJB component method.**

The J2EE specifications restrict the use of RMLTs to single EJB methods. This restriction is because the specifications have no scoping device, beyond a container-imposed method boundary, to which an RMLT can be extended. In WebSphere Application Server Enterprise, you can exploit the **Local Transactions - Boundary** extended deployment setting to give the following advantages:

- Significantly extend the use-cases of RMLTs
- Make conversational interactions with one-phase resource managers possible through ActivitySession support.

An ActivitySession is a WebSphere Application Server Enterprise programming model extension that provides a distributed context with a boundary that is longer than a single method. You can extend the use of RMLTs over the longer ActivitySession boundary, which can be controlled by a client. The ActivitySession boundary reduces the need to use

distributed transactions where ACID operations on multiple resources are not needed. This benefit is exploited through the **Local Transactions - Boundary** extended deployment setting of `ActivitySession`. Such extended RMLTs can remain under the control of the application or be managed by the container depending on the use of the **Local Transactions - Resolution-control** deployment descriptor setting.

#### **Coordinate multiple one-phase resource managers.**

For resource managers that do not support XA transaction coordination, a client can exploit `ActivitySession`-bounded local transaction contexts. Such contexts give a client the same ability to control the completion direction of the resource updates by the resource managers as the client has for transactional resource managers. A client can start an `ActivitySession` and call its entity beans under that context. Those beans can perform their RMLTs within the scope of that `ActivitySession` and return without completing the RMLTs. The client can later complete the `ActivitySession` in a commit or rollback direction and cause the container to drive the `ActivitySession`-bounded RMLTs in that coordinated direction.

To determine how best to configure the transaction support for an application, depending on what you want to do with transactions, consider the following points.

#### **General points**

- You want to start and end global transactions explicitly in the application (BMT session beans and servlets only).  
For a session bean, set the **Transaction type** to `Bean` (to use bean-managed transactions) in the component's deployment descriptor. (You do not need to do this for servlets.)
- You want to access only one XA or non-XA resource in a method.  
In the component's deployment descriptor, set **Local Transactions - Resolution-control** to `ContainerAtBoundary`. In the Container transaction deployment descriptor, set **Transaction** to `Supports`.
- You want to access several XA resources atomically across one or more bean methods.  
In the Container transaction deployment descriptor, set **Transaction** to `Required`, `Requires new`, or `Mandatory`.
- You want to access several non-XA resource in a method without having to worry about managing your own local transactions.  
In the component's deployment descriptor, set **Local Transactions - Resolution-control** to `ContainerAtBoundary`. In the Container transaction deployment descriptor, set **Transaction** to `Not supported`.
- You want to access several non-XA resources in a method and want to manage them independently.  
In the component's deployment descriptor, set **Local Transactions - Resolution-control** to `Application` and set **Local Transactions - Unresolved-action** to `Rollback`. In the Container transaction deployment descriptor, set **Transaction** to `Not supported`.

#### **Points specific to WebSphere Application Server Enterprise**

- You want to access one of more non-XA resources across multiple EJB method calls without having to worry about managing your own local transactions.

In the component's deployment descriptor, set **Local Transactions - Resolution-control** to `ContainerAtBoundary`, **Local Transactions - Boundary** to `ActivitySession`, and **Bean Cache - Activate at** to `ActivitySession`. In the Container transaction deployment descriptor, set **Transaction** to `Not supported` and set **ActivitySession** attribute to `Required`, `Requires new`, or `Mandatory`.

- You want to access several non-XA resources across multiple EJB method calls and want to manage them independently.

In the component's deployment descriptor, set **Local Transactions - Resolution-control** to `Application`, **Local Transactions - Boundary** to `ActivitySession`, and **Bean Cache - Activate at** to `ActivitySession`. In the Container Transaction deployment descriptor, set **Transaction** to `Not supported` and set **ActivitySession** attribute to `Required`, `Requires new`, or `Mandatory`.

- You want to use a single non-XA resource and one or more XAResources.

Use the Last Participant Support of WebSphere Application Server Enterprise.

## Local and global transaction considerations

Applications use resources, such as JDBC data sources or connection factories, that are configured through the Resources view of the WebSphere Application Server Administrative Console. How these resources participate in a global transaction depends on the underlying transaction support of the resource provider. For example, most JDBC providers can provide either XA or non-XA versions of a data source. A non-XA data source can support only resource manager local transactions (RMLTs), but an XA data source can support two-phase commit coordination, as well as local transactions.

If an application uses two or more resource providers that support only RMLTs, then atomicity cannot be assured because of the one-phase nature of these resources. To ensure atomic behavior, the application should use resources that support XA coordination and should access them within a global transaction.

If an application uses only one RMLT, the atomic behavior can be guaranteed by the resource manager, which can be accessed under a local transaction containment context.

An application can also access a single resource manager under a global transaction context, even if that resource manager does not support the XA coordination. An application can do this, because WebSphere Application Server performs an "only resource optimization" and interacts with the resource manager under a RMLT. Within a global transaction context, any attempt to use more than one resource provider that supports only RMLTs causes the global transaction to be rolled back.

At any moment, an instance of an enterprise bean can have work outstanding in either a global transaction context or a local transaction containment context, but never both. An instance of an enterprise bean can change from running under one type of context to the other (in either direction), if all outstanding work in the original context is complete. Any violation of this principle causes an exception to be thrown when the enterprise bean tries to start the new context.

---

## Developing components to use transactions

These topics provide information about developing WebSphere application components to use transactions

The way that applications use transactions depends on the type of application component, as follows:

- A session bean can either use container-managed transactions (where the bean delegates management of transactions to the container) or bean-managed transactions (where the bean manages transactions itself).
- Entity beans use container-managed transactions.
- Web components (servlets) use bean-managed transactions.

You configure whether a component uses container- or bean-managed transactions by setting an appropriate value on the Transaction type deployment attribute, as described in "Setting transactional attributes in the deployment descriptor". You can also configure other transactional deployment descriptor attributes.

If you want a session bean to manage its own transactions, you must write the code that explicitly demarcates the boundaries of a transaction as described in "Using bean-managed transactions".

Similarly, if you want a Web component to use transactions, you must write the code that explicitly demarcates the boundaries of a transaction as described in "Using bean-managed transactions".

## Setting transactional attributes in the deployment descriptor

Use this task to configure the transactional deployment descriptor attributes associated with an EJB or Web module, to enable a J2EE application to use transactions.

To set transactional attributes in the deployment descriptor for an application component (enterprise bean or servlet), complete the following steps:

### Steps for this task

1. (Start the Application Assembly Tool).
2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, click **File-> Open** then select the EAR file.
3. In the navigation pane, select the component instance; for example:
  - For a session bean, expand *ejb\_module\_instance*-> **Session beans** then select the bean instance.
  - For a servlet, expand *web\_application*-> **Web Components** then select the servlet instance.

A property dialog notebook for the component is displayed in the property pane.

4. In the property pane, select the Advanced tab.
5. Set the **Transaction type** attribute, which defines the transactional manner in which the container invokes a method.

You can set this attribute to Container or Bean, as follows:

- For a session bean to use container-managed transactions, set Container



- For a session bean to use bean-managed transactions, set Bean
  - For an entity bean, set Container
  - For a Web component (servlet), set Bean
6. In the property pane, select the IBM Extensions tab.
  7. Configure J2EE component extensions attributes for extended local transaction containment.

To enable management of local transaction containments, configure the following EJB extensions attributes. These attributes configure, for the component, the behaviour of the container's local transaction containment (LTC) environment that the container establishes whenever a global transaction is not present.

#### **Boundary**

Specifies the duration of a local transaction context. You can set this attribute to **Bean method** or **ActivitySession**, as described in ("Entity bean assembly settings").

**Note:** The ActivitySession option is not supported in the web container.

This property can be changed on WAS Enterprise only.

#### **Resolution control**

Specifies how the local transaction is to be resolved before the local transaction context ends: by the application through user code or by the EJB container. You can set this attribute to either **Application** or **ContainerAtBoundary**, as described in ("Entity bean assembly settings").

#### **Unresolved action**

Specifies the action that the container must take when the local transaction context scope ends, if resources are uncommitted by an application in a local transaction and the **Resolution control** is set to Application. You can set this attribute to either **Commit** or **Rollback**, as described in ("Entity bean assembly settings").

8. [For EJB components only] For container-managed transactions, configure how the container must manage the transaction boundaries when delegating a method invocation to an enterprise bean's business method:
  - a. In the navigation pane, select **Container Transactions**.  
This displays a table of the methods for enterprise beans.
  - b. For each method of the enterprise bean set the **Transaction attribute** attribute to an appropriate value, as defined in ("Container transaction assembly settings").

If the application uses ActivitySessions, how the container manages transaction boundaries when delegating a method invocation depends on both the **Transaction attribute** attribute, set here, and the **ActivitySession kind** attribute, as described in "Configuring ActivitySession deployment attributes for an enterprise bean" on page 1033Configuring ActivitySession deployment attributes for an enterprise JavaBean. For more detail about the relationship between these two properties, see "Combining transaction and ActivitySession container policies" on page 1024Combining transaction and ActivitySession container policies.

## Using bean-managed transactions

This topic describes how to enable a session bean or servlet to use bean-managed transactions, to manage its own transactions directly instead of letting the container manage the transactions.

**Note:** Entity beans cannot manage transactions (so cannot use bean-managed transactions).

To enable a session bean or servlet to use bean-managed transactions, complete the following steps:

### Steps for this task

1. Set the **Transaction type** attribute in the component's deployment descriptor to Bean, as described in Setting transactional attributes in the deployment descriptor.
2. Write the component code to actively manage transactions

When writing the code required by a component to manage its own transactions, remember the following basic rules:

- An instance of a stateless session bean cannot reuse the same transaction context across multiple methods called by an EJB client.
- An instance of a stateful session bean can reuse the same transaction context across multiple methods called by an EJB client.

The following code extract shows the standard code required to obtain an object encapsulating the transaction context. There are three basic steps involved:

- The component class must set the value of the `javax.ejb.SessionContext` object reference in the `setSessionContext` method.
- A `javax.transaction.UserTransaction` object is created by calling a lookup on `"java:comp/UserTransaction"`.
- The `UserTransaction` object is used to participate in the transaction by calling transaction methods such as `begin` and `commit` as needed. If an enterprise bean begins a transaction, it must also complete that transaction either by invoking the `commit` method or the `rollback` method.

```
...
import javax.transaction.*;
...
public class MyStatelessSessionBean implements SessionBean {
private SessionContext mySessionCtx =null;
...
public void setSessionContext (SessionContext ctx)throws EJBException {
mySessionCtx =ctx;
}
...
public float doSomething(long arg1)throws FinderException,EJBException {
    UserTransaction userTran = (UserTransaction)initCtx.lookup(
        "java:comp/UserTransaction");
    ...
    //User userTran object to call transaction methods
    userTran.begin ();
    //Do transactional work
    ...
    userTran.commit ();
    ...
}
...
}
```

---

## Configuring transaction properties for an application server

Use this task to configure the transaction properties for an application server; for example, to define the location of the directory that contains the transaction log or to change default timeouts associated with transactions.

To configure the transaction properties for an application server, complete the following steps:

### Steps for this task

1. Start the Administrative console
2. In the navigation pane, select **Servers-> Manage Local Server**  
This displays the properties of the application server in the content pane.
3. Select the Transaction Service tab, to display the properties page for the transaction service, as two notebook pages:

#### **Configuration**

The values of properties defined in the configuration file. If you change these properties, the new values are applied when the application server next starts.

#### **Runtime**

The runtime values of properties. If you change these properties, the new values are applied immediately, but are overwritten with the Configuration values when the application server next starts.

4. Select the Configuration tab, to display the transaction-related configuration properties.
5. **(Optional)** If you want to change the directory in which transaction logs are written, type the full pathname of the directory in the **Transaction log directory** field.

You can check the current runtime value of **Transaction log directory**, by clicking the Runtime tab.

You can also specify a size for the transaction logs, as described in the following step.

**Note:** If you change the transaction log directory, you should apply the change and restart the application server as soon as possible, to minimize the risk of problems caused that might occur before the application server is restarted. For example, if a problem causes the server to fail (with in-flight transactions), the server next starts with the new log directory and is unable to automatically resolve in-flight transactions that were recorded in the old log directory.

6. **(Optional)** If you want to change the default file size of transaction log files, modify the **Transaction log directory** field to include a file size setting, in the following format:

*directory\_name;file\_size*

Where

- *directory\_name* is the name of the transaction log directory
- *file\_size* is the new default size specified in bytes. The *nK* or *nM* suffix can be used to indicate kilobytes or megabytes. If you do not specify a file size value, the default value of 1M is used.

For example, `c:\tranlogs;2M` indicates the files are to be created with 2M bytes size and stored in the directory `c:\tranlogs`.

In a non-production environment, you can use the transaction log directory value of ;0 to disable transaction logging. (There must be no directory name element before the size element of 0.) You should not disable transaction logging in a production environment, because this prevents recovery after a system failure and, therefore, data integrity cannot be guaranteed.

7. In the **Total transaction lifetime timeout** field, type the number of seconds a transaction can remain inactive before it is ended by the transaction service. A value of 0 (zero) indicates that there is no timeout limit.
8. In the **Client inactivity timeout** field, type the number of seconds after which a client is considered inactive and the transaction service ends any transactions associated with that client. A value of 0 (zero) indicates that there is no timeout limit.
9. Click **OK**.
10. Stop then restart the application server.

If you change the transaction log directory configuration property to an incorrect directory name, the application server will restart but be unable to open the transaction logs. You should change the configuration property to a valid directory name, then restart the application server.

## Transaction service settings

Use this page to modify transaction service settings.

To view this administrative console page, click **Servers > Application Servers > server > Transaction Service**.

### Transaction log directory

Specifies the name of a directory for this server where the transaction service stores log files for recovery.

A blank value in the server configuration is expanded by the transaction log at startup as the directory (install\_root)/tranlog/(server\_name).

When the application running on the WebSphere product accesses more than one resource, the WebSphere product stores transaction information to properly coordinate and manage the distributed transaction. In a higher transaction load, this persistence slows down performance of the application server due to its dependency on the operating system and the underlying storage systems.

To achieve better performance, move the transaction log files to a storage device with more physical disk drives, or preferably RAID disk drives. When the log files are moved to the file systems on the RAID disks, the task of writing data to the physical media is shared across the multiple disk drives. This allows more concurrent access to persist transaction information and faster access to that data from the logs. Depending upon the design of the application and storage subsystem, performance gains can range from 10% to 100%, or even more in some cases.

This change is applicable only to the configuration where the application uses distributed resources or XA transactions, for example, multiple databases and resources are accessed within a single transaction. Consider setting this property when the application server shows one or more of following signs:

- CPU utilization remains low despite an increase in transactions
- Transactions fail with several time outs
- Transaction rollbacks occur with *unable to enlist transaction* exception

- Application server hangs in middle of a run and requires the server to be restarted
- The disk on which an application server is running shows higher utilization

|             |  |
|-------------|--|
| Data type   | String   |
| Default     | Initial value is the %WAS_HOME%\tranlog directory and a default size of 1MB.   |
| Recommended | Create a file system with at least 3-4 disk drives RAIDed together in a RAID-0 configuration. Then, create the transaction log on this file system with the default size. When the server is running under load, check the disk input and output. If disk input and output time is more than 5%, consider adding more physical disks to lower the value. If disk input and output is low, but the server is still high, consider increasing the size of the log files. |

### **Total transaction lifetime timeout**

Specifies the maximum duration, in seconds, for transactions on this application server.

Any transaction that is not requested to complete before this time-out is rolled back. If set to 0, there is no time-out limit.

|           |                    |
|-----------|--------------------|
| Data type | Integer            |
| Units     | Seconds            |
| Default   | 120                |
| Range     | 0 to 2 147 483 647 |

### **Client inactivity timeout**

Specifies the maximum duration, in seconds, between transactional requests from a remote client.

Any period of client inactivity that exceeds this timeout results in the transaction rolling back in this application server. If set to 0, there is no timeout limit.

|           |                    |
|-----------|--------------------|
| Data type | Integer            |
| Units     | Seconds            |
| Default   | 60                 |
| Range     | 0 to 2 147 483 647 |

### **Maximum Transaction Timeout**

Specifies the maximum duration, in seconds, that transactions started by or propagated into this application server are allowed to execute.

|           |                |
|-----------|----------------|
| Data type | Integer        |
| Units     | Seconds        |
| Default   | 300            |
| Range     | 0 to 2 147 040 |

---

## **Managing active transactions**

Use this task to manage transactions that are active on an application server.

You can use this task to display a snapshot of all the transactions currently running on an application server. For each transaction, the following properties are shown: its local ID, global ID, and current status. You can also choose to finish transactions manually.

Under normal circumstances, transactions should run and complete (commit or rollback) automatically, without the need for intervention. However, in some circumstances, you may need to finish a transaction manually. For example, you may want to finish a transaction that has become stuck polling a resource manager that you know will not become available again within the desired timeframe.

**Note:** If you choose to finish a transaction on an application server, it is recorded as having completed in the transaction service logs for that server, so will not be eligible for recovery during server start up. If you finish a transaction, you are responsible for cleaning up any in-doubt transactions on the resource managers affected.

To manage the active transactions for an application server, use the administrative console to complete the following steps:

#### Steps for this task

1. In the navigation pane, select **Servers-> Manage Application Servers**  
This displays a list of application servers in the content pane.
2. In the content pane, click *your\_app\_server*  
This displays the properties of the application server, *your\_app\_server*.
3. In the content pane, click the **Runtime** tab.  
This displays the runtime properties of the application server.
4. In the Additional Properties table, select **Transaction Service**  
This displays the runtime properties of the Transaction Service.
5. Click **Manage Transactions**.  
This displays a snapshot of all the transactions currently running on the server. For each transaction, the following properties are shown: its local ID, global ID, and current status.
6. **(Optional)** If you want to finish one or more transactions, select the checkbox provided on the entry for the transaction, then click **Finish**. Alternatively, to finish all transactions, select the checkbox in the header of the transactions table, then click **Finish**.

---

## Managing transaction logging for optimum server availability

This topic describes some considerations and actions that you can use to manage transaction logging to help ensure that the availability of your application servers is optimized.

The transaction service writes information to the transaction log for every global transaction which involves two or more resources or is distributed across multiple servers. The transaction log is stored on disk and is used by the transaction service for recovery after a system or server crash. The transaction log for each application server consists of multiple files held in a single directory. You can change the directory that an application server uses to store the transaction log, as described in "Configuring transaction properties for an application server".

When a global transaction is completed, the information in the transaction log is not needed anymore so is marked for deletion. Periodically, this redundant information is garbage collected and the space reused by new transactions. The log files are created of fixed size at server startup, thus no further disk space allocation is required during the lifetime of the server. The default allocation is suitable for around 500 concurrent transactions.

If all the log space is in use when a transaction needs to save information, the transaction is rolled back and the message "WTRN0075W: The transaction log file is full. Transaction rolled back." is reported to the system error log. No more transactions can commit until more log space is made available when existing active transactions complete.

You can monitor the number of concurrent global transactions by using the performance monitoring counters for transactions. The "Global transaction commit time" counter is a measure of how long a transaction takes to complete and, therefore, how long the log is in use by a transaction. If this value is high, then transactions are taking a long time to complete, which can be due to resource manager or network failures. If you ensure this value is low, the log is more efficiently used and unlikely to become full.

You can change the default size of log files by updating the transaction log settings as described in "Configuring transaction properties for an application server". Take care if you increase the size above the default 1Mbyte setting, because this extends the time of the log file garbage collection process, and can lead to undesirable periodic "transaction stall" effects.

## Configuring transaction aspects of servers for optimum availability

This topic describes some considerations and actions that you can take to configure transaction-related aspects of application servers for optimum availability.

To configure transaction-related aspects of application servers for optimum availability, complete the following steps:

### Steps for this task

1. Store the transaction log files on a fast disk in a highly-available file system, such as a RAID device.

The transaction log may need to be accessed by every global transaction and be used for transaction recovery after a crash. Therefore, the disk the log files are being written to should be on a highly-available file system, such as a RAID device.

The performance of the disk also directly affects the transaction performance. In general, a global transaction makes two disk writes, one after the prepare phase when the outcome of the transaction is known (this information is forced to disk) and a further disk write at transaction completion. Therefore, the transaction logs should be placed on the fastest disks available and not make use of network mounted devices.

2. Mirror the transaction log files by using hardware disk mirroring or dual-ported disks.

If log files have been mirrored or can be recovered, they can be used when restarting a failed server or moved to another machine and another server started there to perform recovery.

Hardware disk mirroring or dual-ported disks can be used by specifying the appropriate file system directory for the transaction logs using the WebSphere Administrative Console.

3. Specify the optimum location of the transaction log directory for application servers.

The default transaction log directory for an application server configuration is unset. By default the application server places transaction log files in a subdirectory of the installed WebSphere tranlog directory (as defined by the WebSphere variable `TRANLOG_ROOT`), and the subdirectory name is the same as the server name. For example, the default directory for a server named `server1` on Windows 2000 is: `c:\WebSphere\AppServer\tranlog\server1`.

You can specify an optimum location of the transaction log directory for all application servers, either on a node or cell scope, by setting the WebSphere variable `TRANLOG_ROOT`. You can specify a different location separately for each application server by setting the **Transaction Log Directory** property for the server.

4. Never allow more than one application server to concurrently use the same set of log files.

Because the transaction logs record the state of global transactions within a server, if the logs become lost or corrupt, then transactions that are in the prepared state before failure can leave resources in an in-doubt state and prevent further updates or access to the resources by other users or servers. These transactions may need to be manually resolved by either committing or rolling back the transactions at the affected resource managers. The failed server can then be cold-started, which creates new empty transaction logs.

If log files have been mirrored or can be recovered, they can be used when restarting the failed server or moved to an alternate server or machine and another server restarted to perform recovery, as described in the related tasks.

Never allow more than one application server to concurrently use the same set of log files, because each server will destroy the information recorded by the other, resulting in corrupt log files that are unusable for future recovery purposes.

5. Configure application servers to always use the same listening port address at each startup.

If you are running distributed transactions between multiple application servers, the remote object references saved in the transaction log need to be redirected to the originating server on recovery.

On Application Server Network Deployment, the node agents automatically redirect such remote object references to the appropriate application servers on recovery. However, if the distributed transaction is between application servers that are not on Application Server Network Deployment, then you must handle the redirection of remote object references for transaction recovery to complete. For example, you must do this if an application server is deployed on WebSphere Application Server (not the Network Deployment edition) and runs distributed transactions with non-WebSphere EJB or Corba servers.

In particular, the default restart action of an application server not on Application Server Network Deployment is to use a different listening port address to the port when the server shut down. This prevents transaction recovery completing. To overcome this, you should always configure application servers to always use the same listening port address at each startup (see the ORB property `com.ibm.CORBA.ListenerPort` in ("Object Request Broker service settings that can be added to the administrative



console”)). You may need to make similar configuration changes to other application servers involved in transactions, to be able to access those servers during recovery.

## Moving a transaction log from one server to another

This topic describes some considerations and actions that you can take to move the transaction logs for an application server to another server.

To move transaction logs from one application server to another, consider the following steps:

### Steps for this task

1. Move all the transaction log files for the application server.

The transaction log directory for each server contains four log files; named tranlog1, tranlog2, XAresource1, and XAresource2. When moving transaction logs from one server to another you must move all four files; otherwise recovery may not complete resulting in data inconsistency.

2. For a single server configuration, move the transaction logs to any server that has access to the same resource managers.

For a single server configuration (where there are no distributed transactions), the transaction logs can be moved to any server (on any node) that has access to the same resource managers as the original server. For example, the server needs communication and valid security access to databases or message queues.

All the transaction log files for the original server need to be moved to a directory accessible by the new server. This can be accomplished by either renaming the transaction log directory or copying all the log files to the new server’s transaction log directory before starting the new server.

**Note:** To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the application server to continue to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

3. For a network-deployed server configuration, move the transaction logs to a server that has the same name and host IP address, and access to the same resource managers.

For a network-deployed server configuration, where there can be distributed transactions present in the logs, there are more restrictions. Distributed transactions that access multiple servers log information about each server involved in the transaction. This information includes the server name and the IP address of the machine on which the server is running. When recovery is taking place on server restart, the server uses this information to contact the distributed servers and similarly, the distributed servers try to contact the server with the same original name. So, if a server fails and the logs need to be recovered on an alternative server, that alternative server needs to have the same name and host IP address as the original server. The alternative server also needs to have access to the same resource managers as the original server. For example, the server needs communication and valid security access to databases or message queues.

**Note:** All servers within a cell must have unique names.

**Note:** To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the

application server to continue to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

## Restarting an application server on a different host

This topic describes some considerations and actions that you can take with transaction logs to restart an application server on a different host.

Moving transactions logs to a different host is similar to moving logs from one server to another, as described in "Moving a transaction log from one server to another".

This involves moving an original application server on one host to an alternative server, which has access to the same resource managers, on another host. For a network-deployed server configuration, the alternative server must have the same name and host IP address as the original server.

**Note:** To complete transaction recovery, the application server uses the resource manager configuration information in the transaction logs. However, for the application server to continue to do new work with the same resource managers, the server must have an appropriate resource manager configuration (as for the original server).

To restart an application server on a different host, complete the following steps:

### Steps for this task

1. Ensure that the alternative application server is stopped.
2. Move all the transaction logs for the original server to the alternative application server, according to the considerations described in "Moving a transaction log from one server to another".
3. Restart the alternative application server.

---

## Transactional interoperation with non-WebSphere application servers

To interoperate transactionally with a non-WebSphere application server, WebSphere Application Server switches dynamically between native transaction contexts and interoperable OTS contexts depending on the capability of the partner with which it is interoperating. The following system properties (that were needed to be set in WebSphere Application Server before version 5.0 to enable transactional interoperation), and the use of native contexts, are deprecated:

```
com.ibm.ejs.jts.jts.ControlSet.nativeOnly=false  
com.ibm.ejs.jts.jts.ControlSet.interoperabilityOnly=true
```

In a future release of WebSphere Application Server only interoperable OTS contexts will be supported.

---

## Troubleshooting transactions

Use this overview task to help resolve a problem that you think is related to the Transaction service.

To identify and resolve transaction-related problems, you can use the standard WebSphere Application Server RAS facilities. If you encounter a problem that you think might be related to transactions, complete the following stages:

### Steps for this task

1. Check for transaction messages in the admin console.  
The Transaction service produces diagnostic messages prefixed by "WTRN". The error message indicates the nature of the problem and provides some detail. The associated message information provides an explanation and any user actions to resolve the problem.
2. Check for Transaction messages in the activity log.  
Activity log messages produced by the Transaction service are accompanied by log analyzer descriptions.
3. Check for more messages in the application server's stdout.log.  
For more information about a problem, check the stdout.log file for the application server, which should contain more error messages and extra details about the problem.
4. Check for messages in the application server's transaction log directory for information about the transactions in-flight when the problem occurred.  
**Note:** If you changed the transaction log directory and a problem caused the application server to fail (with in-flight transactions) before the server was restarted properly, the server will next start with the new log directory and be unable to automatically resolve in-flight transactions that were recorded in the old log directory. To resolve this, you can copy the transaction logs to the new directory then stop and restart the application server.

---

## Transaction service exceptions

This topic lists the exceptions that can be thrown by the WebSphere Application Server transaction service. The exceptions are listed in the following groups:

- Standard exceptions
- Heuristic exceptions

If the EJB container catches a system exception from the business method of an enterprise bean, and the method is running within a container-managed transaction, the container rolls back the transaction before passing the exception on to the client. For more information about how the container handles the exceptions thrown by the business methods for beans with container-managed transaction demarcation, see the section *Exception handling* in the Enterprise JavaBeans 2.0 specification (<http://java.sun.com/products/ejb/2.0.html>). That section specifies the container's action as a function of the condition under which the business method executes and the exception thrown by the business method. It also illustrates the exception that the client receives and how the client can recover from the exception.

### Standard exceptions

The standard exceptions such as `TransactionRequiredException`, `TransactionRolledbackException`, and `InvalidTransactionException` are defined in the Java Transaction API (JTA) 1.0.1 Specification (<http://java.sun.com/products/jta/>).

### **InvalidTransactionException**

This exception indicates that the request carried an invalid transaction context.

**TransactionRequiredException exception**

This exception indicates that a request carried a null transaction context, but the target object requires an active transaction.

**TransactionRolledbackException exception**

This exception indicates that the transaction associated with processing of the request has been rolled back, or marked for roll back. Thus the requested operation either could not be performed or was not performed because further computation on behalf of the transaction would be fruitless.

**Heuristic exceptions**

A heuristic decision is a unilateral decision made by one or more participants in a transaction to commit or rollback updates without first obtaining the consensus outcome determined by the Transaction Service. Heuristic decisions are an issue only after the participant has been prepared and the second phase of commit processing is underway. Heuristic decisions are normally made only in unusual circumstances, such as repeated failures by the transaction manager to communicate with a resource manager during two-phase commit. If a heuristic decision is taken, there is a risk that the decision differs from the consensus outcome, resulting in a loss of data integrity.

The following list provides a summary of the heuristic exceptions. For more detail, see the Java Transaction API (JTA) 1.0.1 Specification (<http://java.sun.com/products/jta/>).

**HeuristicRollback exception**

This exception is raised on the commit operation to report that a heuristic decision was made and that all relevant updates have been rolled back.

**HeuristicMixed exception**

This exception is raised on the commit operation to report that a heuristic decision was made and that some relevant updates have been committed and others have been rolled back.

---

**UserTransaction interface - methods available**

For details about the methods available with the UserTransaction interface, see the WebSphere Application Server application programming interface reference information (Javadoc) or the Java Transaction API (JTA) 1.0.1 Specification (<http://java.sun.com/products/jta/>).

---

**Coordinating access to 1-PC and 2-PC-capable resources within the same transaction**

Use these topics to help you coordinate the use of a single 1-phase commit (1PC) capable resource with any number of 2-phase commit (2PC) capable resources in the same global transaction.

You can coordinate the use of a single 1-phase commit (1PC) capable resource with any number of 2-phase commit (2PC) capable resources in the same global transaction.

At transaction commit, the 2-phase commit resources are prepared first using the 2-phase commit protocol, and if this is successful the 1-phase commit-resource is

then called to `commit(one_phase)`. The 2-phase commit resources are then committed or rolled back depending on the response of the 1-phase commit resource.

For more information about coordinating access to 1PC and 2PC-capable resources within the same transaction, see the following topics:

- Coordinating access to 1-PC and 2-PC-capable resources within the same transaction
- Enabling an application to coordinate access to 1-PC and 2-PC-capable resources within the same transaction
- Configuring an application server to allow logging for heuristic reporting

## Coordinating access to 1-PC and 2-PC-capable resources within the same transaction

You can coordinate the use of a single 1-phase commit (1PC) capable resource with any number of 2-phase commit (2PC) capable resources in the same global transaction.

At transaction commit, the 2-phase commit resources are prepared first using the 2-phase commit protocol, and if this is successful the 1-phase commit-resource is then called to `commit(one_phase)`. The 2-phase commit resources are then committed or rolled back depending on the response of the 1-phase commit resource.

**Note:** If the global transaction is distributed across multiple application servers, you cannot coordinate access to 1-PC and 2-PC-capable resources within the same transaction .

Coordinating access to 1-PC and 2-PC-capable resources within the same transaction introduces an increased risk of an heuristic outcome to the transaction. That is, the transaction manager cannot be sure that all resources were completed in the same direction (either committed or rolled back). For this reason, to enable an application to coordinate access to 1-PC and 2-PC-capable resources within the same transaction, you configure the application to accept the increased risk of an heuristic outcome.

An heuristic outcome occurs if the transaction service (JTS) receives no response from the commit one-phase flow on the 1PC resource. In this situation the transaction service cannot determine whether changes for the 1PC resource were committed or rolled back, so cannot drive reliably the correct outcome of the global transaction on the other 2PC resources.

You can configure the transaction service for an application server to indicate whether or not to log that it is about to commit the 1PC resource. This does not reduce the heuristic hazard, but ensures that any failure, and subsequent recovery, of the application server during the 1PC phase occurs with knowledge of whether or not the 1PC resource was asked to commit:

- If the 1PC resource was asked to commit, a heuristic outcome is reported to the activity log.
- If the 1PC resource was not asked to commit, then the transaction is rolled back consistently.

## Enabling an application to coordinate access to 1-PC and 2-PC-capable resources within the same transaction

Use this task to enable an application to coordinate access to 1-phase and 2-phase commit capable resources within the same transaction.

To enable an application to coordinate access to 1-phase and 2-phase commit capable resources within the same transaction, you must configure the application to accept the increased risk of an heuristic outcome.

To configure an application to indicate that you accept the increased risk of an heuristic outcome, use the Application Assembly tool to complete the following steps:

### Steps for this task

1. Launch the Application Assembly Tool.
2. Open the application EAR file.
3. In the navigation pane, select the application  
This displays the properties notebook in the property pane.
4. In the property pane, select the WAS Enterprise tab.
5. Select the **Accept heuristic hazard** checkbox.
6. To apply the changes and close the Application Assembly Tool, click **OK**.  
Otherwise, to apply the values but keep the property dialog open for additional edits, click **Apply**.
7. **(Optional)** To see changes reflected in your application, regenerate deployment code and re-install the deployable archive.

### **Last participant support extension settings**

Use this page to configure last participant support extensions.

Last participant support is an extension to the transaction service to allow a single one-phase resource to participate in a two-phase transaction with one or more two-phase resources.

To view this administrative console page, click **Applications > Applications > *application\_name* > Last Participant Support Extension**.

**Accept Heuristic Hazard:** Specifies whether the application accepts the possibility of an heuristic hazard occurring in a two-phase transaction containing a one-phase resource.

|         |   |
|---------|---|
| Default | Cleared   |
| Range   | <b>Selected</b><br>The application accepts the increased risk of an heuristic outcome.        |
|         | <b>Cleared</b><br>The application does not accept the increased risk of an heuristic outcome. |

## Configuring an application server to allow logging for heuristic reporting

To enable an application server to log "about to commit 1PC resource" events from transactions that involve a 1-phase commit resource and 2-phase commit resources, use the Administrative console to complete the following steps:

### Steps for this task

1. Start the Administrative console
2. In the navigation pane, select **Servers-> Manage Application Servers-> your\_app\_server**  
This displays the properties of the application server, *your\_app\_server*, in the content pane.
3. Select the Transaction Service tab, to display the properties page for the transaction service, as two notebook pages:

#### **Configuration**

The values of properties defined in the configuration file. If you change these properties, the new values are applied when the application server next starts.

#### **Runtime**

The runtime values of properties. If you change these properties, the new values are applied immediately, but are overwritten with the Configuration values when the application server next starts.

4. Select the Configuration tab, to display the transaction-related configuration properties.
5. Select the **Enable logging for heuristic reporting** checkbox.
6. Click **OK**.
7. Stop then restart the application server.

## Exceptions thrown for transactions involving both single- and two-phase commit resources

The exceptions that can be thrown by transactions that involve single- and two-phase commit resources are the same as those that can be thrown by transactions involving only two-phase commit resources, and are listed in the WebSphere API reference information (<http://publib7b.boulder.ibm.com/wasinfo1/en/info/ee/javadoc/ee/index.html>).

## Last Participant Support: Resources for learning



Use the following links to find relevant supplemental information about Last Participant Support. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming specifications
- Other

## Programming specifications

-  J2EE Activity Service for Extended Transactions
-  Java Transaction API (JTA) 1.0.1

## Other

-  WebSphere Application Server Enterprise Version 5 Overview: Advanced Transactional Connectivity
-  Listing of PDF files to learn about WebSphere Application Server Version 5 ([http://www-3.ibm.com/software/webservers/appserv/appserv\\_v5.html](http://www-3.ibm.com/software/webservers/appserv/appserv_v5.html))
-  Listing of all IBM WebSphere Application Server Redbooks (<http://publib-b.boulder.ibm.com/Redbooks.nsf/Portals/WebSphere>)
-  Listing of all IBM WebSphere Application Server Whitepapers (<http://www-4.ibm.com/software/webservers/appserv/whitepapers.html>)
-  WebSphere Application Server Enterprise Edition 4.0: A Programmer's Guide (<http://www.redbooks.ibm.com/redbooks/SG246504.html>)



---

## Chapter 23. Using the ActivitySession service

These topics provide information about implementing WebSphere enterprise applications that use ActivitySessions.

The ActivitySession service provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. ActivitySessions provide a scoping mechanism for units of work, and both an ActivitySession and a transaction has the same following characteristics:

- It can be bean-managed or container-managed
- It can be distributed across application servers
- It can be used as the context for managing EJB activation policy and lifecycle

An ActivitySession differs significantly from a transaction in the manner of its interaction with resource managers. An ActivitySession is used to scope or coordinate local transactions. That is, an ActivitySession can be used to request multiple 1-phase resource managers to come to an application- or container-determined outcome. Unlike a transaction, an ActivitySession has no notion of a prepare phase or any notion of recovery at a service level.

The WebSphere EJB container and deployment tools support ActivitySessions as an extension to the J2EE programming model. Enterprise beans can be deployed with lifecycles that are influenced by ActivitySession context, as an alternative to transaction context. An enterprise bean with an ActivitySession-scoped lifecycle can participate in a resource manager local transaction (RMLT) that has a duration of the ActivitySession rather than an individual method on the bean (which is all that is possible under the standard J2EE model). Applications can then be composed of several enterprise beans with ActivitySession-based activation, with each bean participating in extended local transactions with one or more resource managers. At the end of the ActivitySession each of the local transactions can be directed to a common outcome by the ActivitySession manager.

You can configure the WebSphere containers and deployable applications to support enterprise beans that operate under application- or container-initiated ActivitySessions rather than, or in addition to, transactions.

For more information about implementing WebSphere enterprise applications that use ActivitySessions, see the following topics:

- The ActivitySession service
  - ActivitySessions and transaction contexts
  - Using ActivitySessions with HTTP sessions
- The ActivitySession service programming interfaces
- Developing a J2EE application to use an ActivitySession
- "Samples: ActivitySessions"
- "Configuring ActivitySession deployment attributes for an enterprise bean"
- "Configuring ActivitySession deployment attributes for a Web application"
- "Disabling or enabling the ActivitySession service"
- Configuring the default ActivitySession timeout
- Troubleshooting ActivitySessions

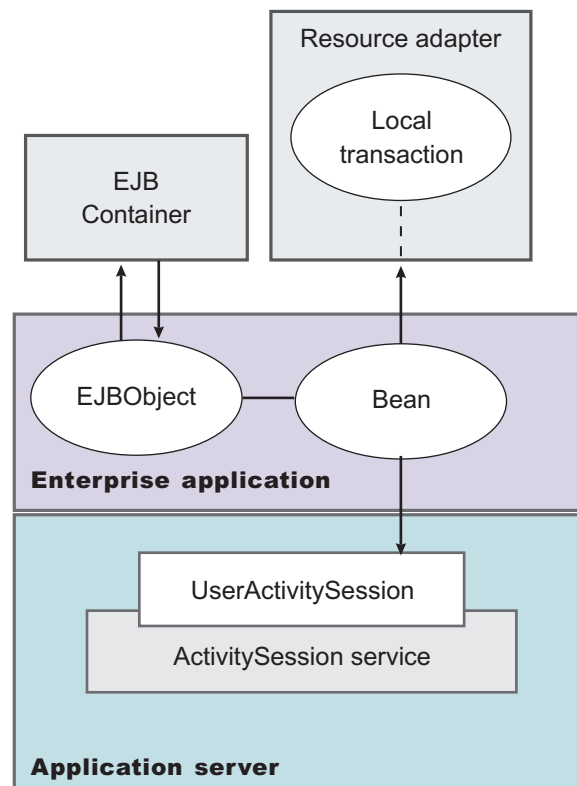
---

## The ActivitySession service

The ActivitySession service provides an alternative unit-of-work (UOW) scope to that provided by global transaction contexts. An ActivitySession context can be longer-lived than a global transaction context and can encapsulate global transactions.

Support for the ActivitySession service is shown in the following figure:

**The ActivitySession service.** This figure shows the main components of the ActivitySession service within WebSphere Application server. For an overview of these components, see the text that accompanies this figure.



Although the purpose of a global transaction is to coordinate multiple resource managers, global transaction context is often used by J2EE applications as a "session" context through which to access EJB instances. An ActivitySession context is such a session context, and can be used in preference to a global transaction in cases where coordination of two-phase commit resource managers is not needed. Further, an ActivitySession can be associated with an HttpSession to extend a "client session" to an HTTP client.

ActivitySession support is available to Web, EJB, and J2EE-client components. EJB components can be divided into beans that exploit container-managed ActivitySessions and beans that use bean-managed ActivitySessions.

The ActivitySession service provides a UserActivitySession application programming interface available to J2EE components that use bean-managed ActivitySessions for application-managed demarcation of ActivitySession context. The ActivitySession service also provides a system programming interface for

container-managed demarcation of `ActivitySession` context and for container-managed enlistment of one-phase resources (RMLTs) in such contexts.

The `UserActivitySession` interface is obtained by a JNDI lookup of `java:comp/websphere/UserActivitySession`. This interface is not available to enterprise beans that use container-managed `ActivitySessions`, and any attempt by such beans to obtain the interface results in a `NotFound` exceptions.

## Using `ActivitySessions` with HTTP sessions

A web application that runs in the WebSphere Web container can participate in an `ActivitySession` context.

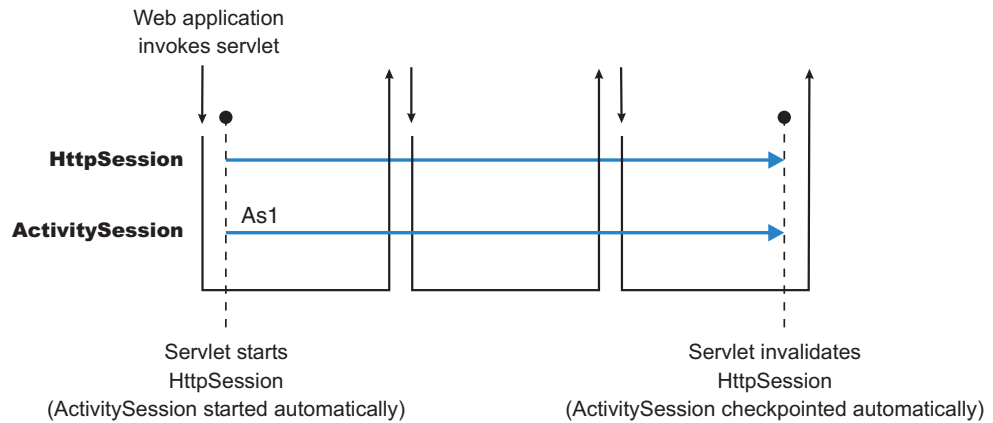
If the web application is designed such that several servlet invocations occur as part of the same logical application, then the servlets can use the `HttpSession` to preserve state across servlet invocations. The `ActivitySession` context is one state that can be suspended into the `HttpSession` and resumed on a future invocation of a servlet that accesses the `HttpSession`.

An `ActivitySession` is associated automatically with an `HttpSession`, so can be used to extend access to the `ActivitySession` over multiple HTTP invocations, over inclusion or forwarding of servlets, and to support EJB activation periods that can be determined by the lifecycle of the web HTTP client. An `ActivitySession` context stored in an `HttpSession` can also be used to relate work for the `ActivitySession` back to a specific web HTTP client.

The Web container manages `ActivitySessions` based on deployment descriptor attributes associated with servlets in the Web application module, as described in "Configuring `ActivitySession` deployment attributes for a Web application". The two usage models are:

- The Web container starts and ends `ActivitySessions`.  
The Web application invokes a servlet that has been configured for container control of `ActivitySessions`.
  - If an `HttpSession` exists then it has an associated `ActivitySession`.
  - If an `HttpSession` does not exist, the servlet can start an `HttpSession`, which causes an `ActivitySession` to be started automatically and associated with the `HttpSession`.

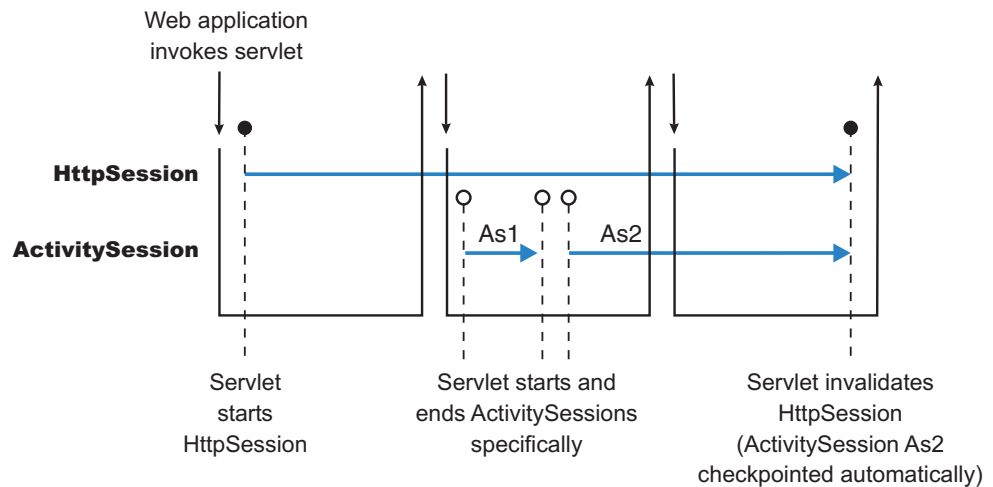
A servlet cannot start a new `HttpSession` until an existing `HttpSession` has been ended. Within an `HttpSession`, the Web application can invoke other servlets that can use the associated `ActivitySession` context. When the Web application invokes a servlet that ends the `HttpSession`, the `ActivitySession` is ended automatically. This is shown in the following diagram:



- The Web application starts and ends ActivitySessions. The Web application invokes a servlet that has been configured for application control of ActivitySessions.
  - If an HttpSession exists and has an associated ActivitySession, the servlet can use or end that ActivitySession context.
  - If an HttpSession does not exist, the servlet can start an HttpSession, but this does not automatically start an ActivitySession.
  - If an HttpSession exists but does not have an associated ActivitySession, the servlet can start a new ActivitySession. This automatically associates the ActivitySession with the HttpSession. The ActivitySession lasts either until the ActivitySession is specifically ended or until the HttpSession is ended.

The servlet cannot start a new ActivitySession until an existing ActivitySession has been ended. The servlet cannot start a new HttpSession until an existing HttpSession has been ended.

Within an HttpSession, the Web application can invoke other servlets that can use or end an existing ActivitySession context or, if no ActivitySession exists start a new ActivitySession. When the Web application invokes a servlet that ends the HttpSession, the ActivitySession is ended automatically. This is shown in the following diagram:



A Web application can invoke servlets configured for either usage model.

The following points apply to both usage models:

- To end an `HttpSession` (and any associated `ActivitySession`), the Web application must invalidate that session. This causes the `ActivitySession` to be checkpointed.
- Any downstream EJBs activated within the context of an `ActivitySession` can be held in memory rather than passivated between servlet invocations, because the client effectively becomes the web HTTP client.
- Web applications can be composed of many Web components, but each component in the Web application must be configured with the same value for `ActivitySessionControl`. `ActivitySessionControl` determines whether the Web component or its container starts any `ActivitySessions`.
- An `ActivitySession` context that encapsulates an active transaction context cannot be associated with an `HttpSession`, because a transaction can hold database locks and should be designed to be shortlived. If an application moves an active transaction to an `HttpSession`, the transaction is rolled back and the `ActivitySession` is suspended into the `HttpSession`. In general, you should design applications to use `ActivitySessions` or other constructs as the long-lived entities and ACID transactions as short-duration entities within these.
- Only one `ActivitySession` can be associated with an `HttpSession` at any time, for the duration of the `ActivitySession`. An `ActivitySession` associated with an `HttpSession` remains associated for the duration of that `ActivitySession`, and cannot be replaced with another until the first `ActivitySession` is completed. The `ActivitySession` can be accessed by multiple servlets if they have shared access to the `HttpSession`.
- `ActivitySessions` are not persistent. If a persistent `HttpSession` exists longer than the server hosting it, any cached `ActivitySession` is terminated when the hosting server ends.
- If the `HttpSession` times out before the associated `ActivitySession` has ended, the `ActivitySession` is reset.
- If the `ActivitySession` times out, it is reset then the `HttpSession` is ended.

## ActivitySession and transaction contexts

The `ActivitySession` service defines a hierarchical relationship between transaction and `ActivitySession` context, requiring that any transaction context be either wholly inside or wholly outside an `ActivitySession` context.

An `ActivitySession` context is very similar to a transaction context and extends the lifecycle choices for activation of enterprise beans; it can encapsulate one or more transactions. The `ActivitySession` context is a distributed context that, like the transaction context, can be bean- or container-managed. An `ActivitySession` context is used mainly by a client to scope the lifecycle of an enterprise bean that it uses either beyond or in the absence of individual transactions started by that client.

`ActivitySessions` have a lower overhead than transactions and can be used instead of transactions that are only used to scope the lifecycle of a called enterprise bean. For a bean with an activation policy of `ActivitySession`, the duration of any resource manager local transactions (RMLTs) started by that bean can be bounded by the duration of the `ActivitySession` instead of the bean method in which the RMLT was started. This provides flexibility and potential for using RMLTs in an enterprise bean beyond the scenarios described in the J2EE specifications. The J2EE specifications define that RMLTs need to be completed before the end of the bean method, because the bean method is the only containment boundary for local transactions available in those specifications.

The following rules defines the relationship between transactions and ActivitySessions.

- The EJB or Web container always uses a local transaction containment (LTC) if there is no global transaction present. An LTC can be method-scoped or ActivitySession-scoped.
- Before a method dispatch, the container ensures that there is always either an LTC or global transaction context, but never both contexts.
- ActivitySessions cannot be nested within each other. Any attempt to start a nested ActivitySession results in a `com.ibm.websphere.ActivitySession.NotSupportedException` on `UserActivitySession.beginSession()`.
- An ActivitySession can wholly encapsulate one or more global transactions.
- An ActivitySession cannot be encapsulated by a global transaction nor should ActivitySession and global transaction boundaries overlap. Any attempt to start an ActivitySession in the presence of a global transaction context results in a `com.ibm.websphere.ActivitySession.NotSupportedException` on `UserActivitySession.beginSession()`. Any attempt to call `endSession(EndModeCheckpoint)` on an ActivitySession that contains an incomplete global transaction results in a `com.ibm.websphere.ActivitySession.ContextPendingException`. Neither the global transaction nor the ActivitySession context are affected. If `endSession(EndModeReset)` is called then the ActivitySession is reset and the global transactions marked `rollback_only`.
- Each global transaction wholly encapsulated by an ActivitySession is independent of every other global transaction within that ActivitySession. A rollback of one global transaction does not affect any others or the ActivitySession itself.
- ActivitySession and global transaction contexts can coexist with an ActivitySession encapsulating one or more serially-executing global transactions.

## Combining transaction and ActivitySession container policies

This topic provides details about the relationship between the deployment descriptor properties that determine how the container manages ActivitySession boundaries.

If an enterprise bean uses ActivitySessions, how the EJB container manages ActivitySession boundaries when delegating a method invocation depends on both the **ActivitySession kind** and **Transaction attribute** deployment descriptor attributes configured for the enterprise bean. The following table lists the relationship between these two properties.

In each row, the final column describes the behavior that the EJB container takes with respect to global transaction and ActivitySession context, based on the following abbreviations:

***Sn*** An ActivitySession, where *n* indicates the ActivitySession instance.

***Tn*** A transaction, where *n* indicates the transaction instance.

In every case where the container does not start or leave a global transaction context associated with the thread, it starts (or obtains from the bean instance) a local transaction containment and associates that with the thread. The duration of the local transaction containment is determined by a combination of the local-transaction boundary descriptor (configured as part of the application

deployment descriptor, and not shown in the following table) and the presence or not of an ActivitySession context, as described in "ActivitySession and transaction contexts".

The rows highlighted in bold are not allowed.

**Container behavior for activitysession and transaction policies deployment settings**

| Bean ActivitySession policy(ActivitySession kind) | Bean transaction policy(Transaction attribute) | Received contexts | Container behaviour            |                                |
|---|--|-------------------|--------------------------------|--------------------------------|
| Required  | Required                                       | None              | Start S1, Start T1             |                                |
|   |  | S1                | Start T1                       |                                |
|   |  | T1                | Suspend T1, Start S1, Start T2 |                                |
|   | Requires new                                   | Requires new      | S1, T1                         | No Action                      |
|   |  |                   | None                           | Start S1, Start T1             |
|   |  |                   | S1                             | Start T1                       |
|   | Supports                                       | Supports          | T1                             | Suspend T1, Start S1, Start T2 |
|   |  |                   | S1, T1                         | Suspend T1, Start T2           |
|   |  |                   | None                           | Start S1                       |
|   | Not supported                                  | Not supported     | S1                             | No Action                      |
|   |  |                   | T1                             | Suspend T1, Start S1           |
|   |  |                   | None                           | Start S1                       |
|   | Mandatory                                      | Mandatory         | S1                             | No Action                      |
|   |  |                   | T1                             | Suspend T1, Start S1           |
|   |  |                   | None                           | Exception                      |
|   | Never  | Never             | S1                             | Exception                      |
|   |  |                   | T1                             | Exception                      |
|   |  |                   | None                           | No action                      |
|   |  | S1, T1            | No action                      |                                |
|   |  | None              | Start S1                       |                                |
|   |  | S1                | No Action                      |                                |
|   |  | T1                | Suspend T1, Start S1           |                                |
|   |  | S1, T1            | Exception                      |                                |
|   |  | None              | Exception                      |                                |

| Bean ActivitySession policy(ActivitySession kind) | Bean transaction policy(Transaction attribute) | Received contexts         | Container behaviour            |
|---|--|---------------------------|--------------------------------|
| Requires new                                      | Required                                       | None                      | Start S1 + T1                  |
|   |  | S1                        | Suspend S1, Start S2 + T1      |
|   |  | T1                        | Suspend T1, Start S1 + T2      |
|   |  | S1 + T1                   | Suspend S1 + T1, Start S2 + T2 |
|   | Requires new                                   | None                      | Start S1 + T1                  |
|   |  | S1                        | Suspend S1, Start S2 + T1      |
|   |  | T1                        | Suspend T1, Start S1 + T2      |
|   |  | S1 + T1                   | Suspend S1 + T1, Start S2 + T2 |
|   | Supports                                       | None                      | Start S1                       |
|   |  | S1                        | Suspend S1, Start S2           |
|   |  | T1                        | Suspend T1, Start S1           |
|   |  | S1, T1                    | Suspend S1 + T1, Start S2      |
|   | Not supported                                  | None                      | Start S1                       |
|   |  | S1                        | Suspend S1, Start S2           |
|   |  | T1                        | Suspend T1, Start S1           |
|   |  | S1, T1                    | Suspend S1 + T1, Start S2      |
|   | <b>Mandatory</b>                               | <b>None</b>               | <b>Exception</b>               |
|   |  | <b>S1</b>                 | <b>Exception</b>               |
|   |  | <b>T1</b>                 | <b>Exception</b>               |
|   |  | <b>S1, T1</b>             | <b>Exception</b>               |
| Never   | None   | Start S1                  |                                |
|   | S1   | Suspend S1, Start S2      |                                |
|   | T1   | Suspend T1, Start S1      |                                |
|   | S1, T1   | Suspend S1 + T1, Start S2 |                                |



| Bean ActivitySession policy(ActivitySession kind) | Bean transaction policy(Transaction attribute) | Received contexts | Container behaviour  |
|---|--|-------------------|----------------------|
| Supports  | Required                                       | None              | Start T1             |
|   |  | S1                | Start T1             |
|   |  | T1                | No Action            |
|   |  | S1, T1            | No Action            |
|   | Requires new                                   | None              | Start T1             |
|   |  | S1                | Start T1             |
|   |  | T1                | Suspend T1, Start T2 |
|   |  | S1, T1            | Suspend T1, Start T2 |
|   | Supports                                       | None              | No Action            |
|   |  | S1                | No Action            |
|   |  | T1                | No Action            |
|   |  | S1, T1            | No Action            |
|   | Not supported                                  | None              | No Action            |
|   |  | S1                | No Action            |
|   |  | T1                | Suspend T1           |
|   |  | S1, T1            | Suspend T1           |
|   | Mandatory                                      | None              | Exception            |
|   |  | S1                | Exception            |
|   |  | T1                | No Action            |
|   |  | S1, T1            | No Action            |
|   | Never  | None              | No Action            |
| S1  |  | No Action         |                      |
| T1  |  | Exception         |                      |
| S1, T1  |  | Exception         |                      |

| Bean ActivitySession policy(ActivitySession kind) | Bean transaction policy(Transaction attribute) | Received contexts | Container behaviour       |
|---|--|-------------------|---------------------------|
| Not supported                                     | Required                                       | None              | Start T1                  |
|   |  | S1                | Suspend S1, Start T1      |
|   |  | T1                | No Action                 |
|   |  | S1, T1            | Suspend S1 + T1, Start T2 |
|   | Requires new                                   | None              | Start T1                  |
|   |  | S1                | Suspend S1, Start T1      |
|   |  | T1                | Suspend T1, Start T2      |
|   |  | S1, T1            | Suspend S1 + T1, Start T2 |
|   | Supports                                       | None              | No Action                 |
|   |  | S1                | Suspend S1                |
|   |  | T1                | No Action                 |
|   |  | S1, T1            | Suspend S1 + T1           |
|   | Not supported                                  | None              | No Action                 |
|   |  | S1                | Suspend S1                |
|   |  | T1                | Suspend T1                |
|   |  | S1, T1            | Suspend S1 + T1           |
|   | Mandatory                                      | None              | Exception                 |
|   |  | S1                | Exception                 |
|   |  | T1                | No Action                 |
|   |  | S1,T1             | Exception                 |
|   | Never  | None              | No Action                 |
|   |  | S1                | Suspend S1                |
|   |  | T1                | Exception                 |
|   |  | S1, T1            | Suspend S1 + T1           |

| Bean ActivitySession policy(ActivitySession kind) | Bean transaction policy(Transaction attribute) | Received contexts | Container behaviour |                      |
|---|--|-------------------|---------------------|----------------------|
| Mandatory   | Required                                       | None              | Exception           |                      |
|   |  | S1                | Start T1            |                      |
|   |  | T1                | Exception           |                      |
|   | Requires new                                   | None              | S1, T1              | No Action            |
|   |  |                   | None                | Exception            |
|   |  |                   | S1                  | Start T1             |
|   | Supports                                       | None              | T1                  | Exception            |
|   |  |                   | S1, T1              | No Action            |
|   |  |                   | None                | Exception            |
|   | Not supported                                  | None              | S1, T1              | Suspend T1, Start T2 |
|   |  |                   | S1                  | No Action            |
|   |  |                   | T1                  | Exception            |
|   | Mandatory                                      | None              | S1, T1              | Suspend T1           |
|   |  |                   | S1                  | Exception            |
|   |  |                   | T1                  | Exception            |
|   | Never  | None              | S1, T1              | No Action            |
|   |  |                   | S1                  | No Action            |
|   |  |                   | T1                  | Exception            |
|   | Never  | Required          | S1, T1              | Exception            |
|   |  |                   | None                | Start T1             |
|   |  |                   | S1                  | Exception            |
| Requires new                                      |  | None              | T1                  | No Action            |
|   |  |                   | S1, T1              | Exception            |
|   |  |                   | None                | Start T1             |
| Supports  |  | None              | S1                  | Exception            |
|   |  |                   | T1                  | No Action            |
|   |  |                   | S1, T1              | Exception            |
| Not supported                                     |  | None              | S1, T1              | Exception            |
|   |  |                   | S1                  | No Action            |
|   |  |                   | T1                  | Suspend T1           |
| Mandatory   |  | None              | S1, T1              | Exception            |
|   |  |                   | S1                  | Exception            |
|   |  |                   | T1                  | No Action            |
| Never   |  | None              | S1, T1              | Exception            |
|   |  |                   | S1                  | No Action            |
|   |  |                   | T1                  | Exception            |
|   |  |                   | S1, T1              | Exception            |

| Bean ActivitySession policy(ActivitySession kind) | Bean transaction policy(Transaction attribute) | Received contexts | Container behaviour |
|---|--|-------------------|---------------------|
| Bean managed                                      | Bean managed                                   | None              | No Action           |
|   |  | S1                | Suspend S1          |
|   |  | T1                | Suspend T1          |
|   |  | S1, T1            | Suspend S1 + T1     |

## Developing a J2EE application to use ActivitySessions

This topic provides an overview of the scenarios for which you would develop a J2EE application to use an ActivitySession.

The following common J2EE application scenarios make use of an ActivitySession:

- Developing a J2EE application to use one or more enterprise beans that are persisted to non-transactional datastores.

This scenario can be used by an application that needs to coordinate multiple 1-phase resource managers; for example, for two or more entity EJBs whose persistence is delegated to LocalTransaction resource adapters.

In this scenario, the enterprise beans used by the application have an Activation policy of ActivitySession and a local transaction containment policy with a boundary of ActivitySession and resolution-control of ContainerAtBoundary. These configuration attributes are described in "Configuring ActivitySession deployment attributes for an enterprise bean". The synchronization of the EJB state data is synchronized, by the container, with the 1-phase resource managers at ActivitySession completion and no application code is required to be aware of ActivitySession support.

- Developing a J2EE application in which an enterprise bean accesses a resource manager multiple times in different business methods.

This scenario can be used by an application that needs to extend a resource manager local transaction (RMLT) over several business methods of an enterprise bean instance.

In this scenario, the enterprise beans used by the application have an Activation policy of ActivitySession and a local transaction containment policy with a boundary of ActivitySession and resolution-control of Application. These configuration attributes are described in "Configuring ActivitySession deployment attributes for an enterprise bean". The application logic starts and ends the RMLTs, for example using the javax.resource.cci.LocalTransaction interface offered by a LocalTransaction Connector, but is not constrained to start and commit the LocalTransaction in the same method.

- Developing a J2EE client application to use an ActivitySession to scope EJB activation and load-balancing.

This scenario can be used by an application client that needs to access an entity bean instance several times in the same client session, either without needing to run under a transaction context, or with the need to run under a number of distinct and serially-executed transactions.

In this scenario, the enterprise beans used by the application client have an Activation policy of ActivitySession and a local transaction containment policy appropriate to the function of the enterprise bean. These configuration attributes are described in "Configuring ActivitySession deployment attributes for an enterprise bean". The J2EE client application can represent a period of user activity, for example a signon period, during which a number of interactions

occur with one or more enterprise beans. If the J2EE client application begins an `ActivitySession` and invokes the enterprise beans within the scope of the UOW represented by the `ActivitySession`, then the enterprise bean instances are activated by the container on the `ActivitySession` boundary and remain in the active state until passivated by the container at the end of the `ActivitySession`. Workload affinity management based on the `ActivitySession` is a platform quality of service. Global transactions can begin and end within the `ActivitySession`, if they are wholly encapsulated by the `ActivitySession` and run serially. EJB instances activated at the `ActivitySession` boundary remain active across the serial global transactions.

- Developing a Web application client to participate in an `ActivitySession` context. A web application that runs in the WebSphere Web container can participate in an `ActivitySession` context. Web applications can use the `UserActivitySession` interface to begin and end an `ActivitySession` context. Also, the `ActivitySession` can be associated with an `HttpSession`, thereby extending access to the `ActivitySession` over multiple HTTP invocations and supporting EJB activation periods that can be determined by the lifecycle of the web HTTP client. The Web container manages `ActivitySessions` based on deployment descriptor attributes associated with the Web application module, as described in "Configuring `ActivitySession` deployment attributes for a Web application".

#### **General considerations:**

- An application that is accessed under an `ActivitySession` context can receive a `javax.transaction.InvalidTransactionException RemoteException`, thrown by the EJB container when servicing any application method. This exception occurs when an instance of an enterprise bean that has an `ActivitySession`-based activation policy becomes involved with concurrent global and local transactions.
- To enable an enterprise bean to participate in an `ActivitySession` context and support `ActivitySession`-based operations, it must be configured with an `ActivationPolicy` of `ACTIVITY_SESSION`. A bean configured with `ActivationPolicy` of either `TRANSACTION` or `ONCE` cannot participate in `ActivitySession` context.
- A session bean can either use container-managed `ActivitySessions` or implement bean-managed `ActivitySessions`; entity beans can only use container-managed `ActivitySessions`. A bean is deployed to be bean-managed or container-managed with respect to `ActivitySession` management by setting its transaction type deployment attribute to be bean-managed or container-managed when deploying the enterprise bean, as described in "Configuring `ActivitySession` deployment attributes for an enterprise bean". A bean that uses bean-managed transactions can use bean-managed `ActivitySessions`; a bean that uses container-managed transactions can use container-managed `ActivitySessions`.
- If you want a session bean or J2EE client to manage its own `ActivitySessions`, you must write the code that explicitly demarcates the boundaries of an `ActivitySession`, as described in "Developing an enterprise bean or J2EE client to manage `ActivitySessions`".

For examples of using `ActivitySessions` in J2EE applications, see "Samples: `ActivitySessions`".

---

## Developing an enterprise bean or J2EE client to manage ActivitySessions

Use this task to write the code needed by a session EJB or J2EE client application to manage an ActivitySession, based on the example code extract provided.

In most situations, an enterprise bean can depend on the EJB container to manage ActivitySessions within the bean. In these situations, all you need to do is set the appropriate ActivitySession attributes in the EJB module deployment descriptor, as described in "Configuring ActivitySession deployment attributes for an enterprise bean". Further, in general, it is practical to design your enterprise beans so that all ActivitySession management is handled at the enterprise bean level.

However, in some cases you may need to have a session bean or J2EE client participate directly in ActivitySessions. You then need to write the code needed by the session bean or J2EE client application to manage its own ActivitySessions.

**Note:** Session beans that use BMT and have an **Activate at** setting of Activity session can manage ActivitySessions. Entity beans cannot manage ActivitySessions; the EJB container always manages ActivitySessions within entity beans.

When preparing to write code needed by a session bean or J2EE client application to manage ActivitySessions, consider the points described in "ActivitySession and transaction contexts".

To write the code needed by a session EJB or J2EE client application to manage an ActivitySession, complete the following steps based on the example code extract below:

### Steps for this task

1. Get an initial context for the ActivitySession.
2. Get an implementation of the UserActivitySession interface, by a JNDI lookup of the URL `java:comp/websphere/UserActivitySession`. The UserActivitySession interface is used to begin and end ActivitySessions and to query various attributes of the active ActivitySession associated with the thread.
3. **(Optional)** Set the timeout, in seconds, after which any subsequently started ActivitySessions are automatically completed by the ActivitySession service. If the session bean or J2EE client does not specifically set this value, the default timeout (300 seconds) is used.

The default timeout can also be overridden for each application server, on the *server*-> **Activity Session Service** panel of the administrative console.

4. Start the ActivitySession, by calling the `beginSession()` method of the UserActivitySession.
5. Within the ActivitySession, call business methods to do the work needed. You can also call UserActivitySession methods to manage the ActivitySession; for example, to get the status of the ActivitySession or to checkpoint all the ActivitySession resources involved in the ActivitySession.
6. End the ActivitySession, by calling the `endSession()` method of the UserActivitySession.

### Usage scenario

The following code extract provides a basic example of using the `UserActivitySession` interface:

```
// Get initial context
InitialContext ic = new InitialContext();
// Lookup UserActivitySession
UserActivitySession uas = (UserActivitySession)ic.lookup("java:comp/websphere/UserActivitySession");

// Set the ActivitySession timeout to 60 seconds
uas.setSessionTimeout(60);
// Start a new ActivitySession context
uas.beginSession();
// Do some work under this context
MyBeanA beanA.doSomething();
...
MyBeanB beanB.doSomethingElse();
// End the context
uas.endSession(EndModeCheckpoint);
```

---

## Configuring ActivitySession deployment attributes for an enterprise bean

Use this task to configure the `ActivitySession` deployment attributes for an enterprise bean to enable the bean to participate in an `ActivitySession` context and support `ActivitySession`-based operations.

You can specify `ActivitySession` deployment attributes as part of the deployment of an enterprise bean.

To configure the `ActivitySession` deployment attributes for an enterprise bean, use the Application Assembly Tool to complete the following steps:

### Steps for this task

1. Launch the Application Assembly Tool.
2. Create or edit the application EAR file.  
For example, to change attributes of an existing application, click **File-> Open** then select the EAR file.
3. In the navigation pane, select the enterprise bean instance; for example, for an entity bean expand *ejb\_module\_instance*-> **Entity beans** then select the bean instance.  
A property dialog notebook for the enterprise bean is displayed in the property pane.
4. In the property pane, select the IBM Extensions tab.
5. In the Bean Cache group box, set the **Activate at** attribute to **Activity Session**:  
An enterprise bean with this activation policy is activated and passivated as follows:
  - On an `ActivitySession` boundary, if an `ActivitySession` context is present on activation.
  - On a transaction boundary, if a transaction context, but no `ActivitySession` context, is present on activation.
  - Otherwise on an invocation boundary.
6. **(Optional)** In the Local Transactions group box, set the **Boundary** attribute to **ActivitySession**:

When this setting is used, the local transaction must be resolved within the scope of any `ActivitySession` in which it was started or, if no `ActivitySession` context is present, within the same bean method in which it was started.

7. For entity beans, or session beans, set the `ActivitySessions` properties for each EJB method.
  - a. In the navigation pane, select **Container ActivitySessions**.  
This displays a table of the methods for enterprise beans.
  - b. For each method of the enterprise bean set the **ActivitySession kind** attribute to specify how the container must manage the `ActivitySession` boundaries when delegating a method invocation to an enterprise bean's business method:

#### **Supports**

If the client invokes the bean method within an `ActivitySession`, the container invokes the bean method within an `ActivitySession` context. If the client invokes the bean method without a `ActivitySession` context, the container invokes the bean method without an `ActivitySession` context. The `ActivitySession` context is passed to any enterprise bean objects or resources that are used by this bean method.

#### **Not supported**

The container invokes bean methods without an `ActivitySession` context. If a client invokes a bean method from within an `ActivitySession` context, the container suspends the association between the `ActivitySession` and the current thread before invoking the method on the enterprise bean instance. The container then resumes the suspended association when the method invocation returns. The suspended `ActivitySession` context is not passed to any enterprise bean objects or resources that are used by this bean method.

**Never** The container invokes bean methods without an `ActivitySession` context.

- If the client invokes a bean method from within an `ActivitySession` context, the container throws an `InvalidActivityException` exception, which is a `javax.rmi.RemoteException`.
- If the client invokes a bean method from outside an `ActivitySession` context, the container behaves in the same way as if the **Not Supported** value was set. The client must call the method without an `ActivitySession` context.

#### **Required**

The container invokes the bean method within an `ActivitySession` context. If a client invokes a bean method from within an `ActivitySession` context, the container invokes the bean method within the client `ActivitySession` context. If a client invokes a bean method outside an `ActivitySession` context, the container creates a new `ActivitySession` context and invokes the bean method from within that context. The `ActivitySession` context is passed to any enterprise bean objects or resources that are used by this bean method.

#### **Requires new**

The container always invokes the bean method within a new `ActivitySession` context, regardless of whether the client invokes the



method within or outside an ActivitySession context. The new ActivitySession context is passed to any enterprise bean objects or resources that are used by this bean method.

Any received ActivitySession context is suspended for the duration of the method and resumed after the method ends. The container starts a new ActivitySession before method dispatch and completes it after the method ends.

### **Mandatory**

The container always invokes the bean method within the ActivitySession context associated with the client. If the client attempts to invoke the bean method without an ActivitySession context, the container throws an ActivityRequiredException exception to the client. The ActivitySession context is passed to any EJB object or resource accessed by an enterprise bean method.

The ActivityRequiredException exception is `javax.rmi.RemoteException`.

How the container manages the ActivitySession boundaries when delegating a method invocation depends on both the **ActivitySession kind** attribute, set here, and the **Transaction attribute** attribute, as described in "Setting transactional attributes in the deployment descriptor" (not in this document). For more detail about the relationship between these two properties, see "Combining transaction and ActivitySession container policies".

8. To apply the changes and close the Application Assembly Tool, click **OK**. Otherwise, to apply the values but keep the property dialog open for additional edits, click **Apply**.
9. **(Optional)** To see changes reflected in your application, regenerate deployment code and reinstall the deployable archive.

## **Container ActivitySession assembly properties for EJB modules**

Use this page to specify how a container must manage the scope of an ActivitySession for an enterprise bean's method invocations.

### **Name**

Specifies a name for the mapping between an ActivitySession attribute and one or more methods.

Datatype                      String

### **Description**

Contains text that describes the mapping

Datatype                      String

### **Methods**

The methods to which the ActivitySession attribute applies.

To add a new method, click **Add**. Expand the tree to select the method or methods from the EJB module

## ActivitySession attribute

How the container must manage the activity session boundaries when delegating a method invocation to an enterprise bean's business method

Default Supports

Range

### Not supported

The container invokes bean methods without an ActivitySession context. If a client invokes a bean method from within an ActivitySession context, the container suspends the association between the ActivitySession and the current thread before invoking the method on the enterprise bean instance. The container then resumes the suspended association when the method invocation returns. The suspended ActivitySession context is not passed to any enterprise bean objects or resources that are used by this bean method.

### Required

The container invokes the bean method within an ActivitySession context. If a client invokes a bean method from within an ActivitySession context, the container invokes the bean method within the client ActivitySession context. If a client invokes a bean method outside an ActivitySession context, the container creates a new ActivitySession context and invokes the bean method from within that context. The ActivitySession context is passed to any enterprise bean objects or resources that are used by this bean method.

### Supports

If the client invokes the bean method within an ActivitySession, the container invokes the bean method within an ActivitySession context. If the client invokes the bean method without a ActivitySession context, the container invokes the bean method without an ActivitySession context. The ActivitySession context is passed to any enterprise bean objects or resources that are used by this bean method.

### Requires new

The container always invokes the bean method within a new ActivitySession context, regardless of whether the client invokes the method within or outside an ActivitySession context. The new ActivitySession context is passed to any enterprise bean objects or resources that are used by this bean method.

Any received ActivitySession context is suspended for the duration of the method and resumed after the method ends. The container starts a new ActivitySession before method dispatch and completes it after the method ends.

Default Supports

Range  
(continued)

### **Mandatory**

The container always invokes the bean method within the `ActivitySession` context associated with the client. If the client attempts to invoke the bean method without an `ActivitySession` context, the container throws an `ActivityRequiredException` exception to the client. The `ActivitySession` context is passed to any EJB object or resource accessed by an enterprise bean method.

The `ActivityRequiredException` exception is defined as a `javax.rmi.RemoteException` that is propagated over an ORB boundary as a `CORBA.ACTIVITY_REQUIRED` system exception.

EJB clients that access these entity beans must do so within an existing `ActivitySession`. For other enterprise beans, the enterprise bean or bean method must implement bean-managed `ActivitySessions` or use the **Required** or **Requires New** value. For non-enterprise bean EJB clients, the client must invoke an `ActivitySession` by using the `UserActivitySession` interface.

### **Never**

The container invokes bean methods without an `ActivitySession` context.

- If the client invokes a bean method from within an `ActivitySession` context, the container throws an `InvalidActivityException` exception, which is defined as a `javax.rmi.RemoteException` that is propagated over an ORB boundary as a `CORBA.INVALID_ACTIVITY` system exception.
- If the client invokes a bean method from outside an `ActivitySession` context, the container behaves in the same way as if the **Not Supported** value was set. The client must call the method without an `ActivitySession` context.

---

## Configuring `ActivitySession` deployment attributes for a Web application

Use this task to configure the `ActivitySession` deployment attributes for a Web application to start `UserActivitySessions` and perform work scoped within `ActivitySessions`.

You can specify `ActivitySession` deployment attributes as part of the deployment of a Web application.

To configure the `ActivitySession` deployment attributes for a Web application, use the Application Assembly Tool to complete the following steps:

### Steps for this task

1. Launch the Application Assembly Tool.
2. Create or edit the Web module.  
For example, to change attributes of an existing module, click **File-> Open** then select the archive file for the module.
3. In the navigation pane, expand **web\_application-> Web Components** then select the servlet instance.  
A property dialog box for the servlet instance is displayed in the property pane.
4. In the property pane, select the WAS Enterprise tab.  
This displays the Enterprise properties in the property pane.

5. Set the **ActivitySession control kind** attribute to either Application, Container, or None. All Web components in a Web application must be configured with the same value for ActivitySession control kind.

#### Application

The Web application is responsible for starting and ending ActivitySessions, as follows:

- If an HttpSession is active when an application begins an ActivitySession, then the container associates the ActivitySession with the HttpSession.
- If an ActivitySession is started in the absence of an HttpSession, then the application must ensure it is completed before the dispatched method completes; otherwise, an exception results.
- If an HttpSession is associated with a request dispatched to an application with this ActivitySession control value, and if that HttpSession has an ActivitySession associated with it, then the container dispatches the request in the context of that ActivitySession. For example, the container resumes the ActivitySession context onto the thread before the dispatch.
- A Web application can use both transactions and ActivitySessions. Any transactions started within the scope of an ActivitySession must be ended by the web component that started them and within the same request dispatch.

#### Container

A servlet has no access to UserActivitySessions. Any HttpSession started by the servlet has an ActivitySession automatically associated with it by the container, and this ActivitySession is put onto the thread of execution. If such a servlet is dispatched by a request that has an HttpSession containing no ActivitySession, then the container starts an ActivitySession and associates it with the HttpSession and the thread.

A Web application can use both transactions and ActivitySessions. Any transactions started within the scope of an ActivitySession must be ended by the web component that started them and within the same request dispatch.

**None** A servlet has no access to UserActivitySessions, and no participation in an ActivitySession is tolerated. Any HttpSession containing an ActivitySession that is associated with a request dispatched on such a servlet is rejected with a ServletException.

6. To apply the changes and close the Application Assembly Tool, click **OK**. Otherwise, to apply the values but keep the property dialog open for additional edits, click **Apply**.
7. **(Optional)** To see changes reflected in your application, regenerate deployment code and re-install the deployable archive.

---

## Disabling or enabling the ActivitySession service

Use this task to disable or enable the ActivitySession service for an application server.

You can use the ActivitySession **Startup** property to specify whether or not the ActivitySession service is started automatically for an application server.

To configure the ActivitySession **Startup** property for an application server, use the Administrative console to complete the following steps:

#### Steps for this task

1. Start the Administrative console.
2. In the navigation pane, expand **Servers-> Manage Application Servers**  
This displays a list of the application servers in the content pane.
3. In the Content pane, select the application server that you want to configure.  
This displays the properties for the application server in the content pane.
4. In the Additional Properties table, select **ActivitySession service**. This displays the ActivitySession service properties in the content pane.
5. Select or clear the **Startup** property as needed:

#### **Selected**

[Default] The ActivitySession service is started when the application server is started. This enables applications that specify use of ActivitySessions in their deployment descriptors to run on such an application server.

#### **Cleared**

The ActivitySession service is not started when the application server is started. Applications that specify use of ActivitySessions in their deployment descriptors cannot start on such an application server.

Any attempt to start an application that uses ActivitySessions is rejected and a message issued:

```
WACS0043E: Error found starting an application. application_name
specified an ActivitySession attribute that is not allowed when the
ActivitySession service is not enabled
```

If this happens during server startup, the server continues to start without the application.

6. Click **OK**.
7. To save your configuration, click **Save** on the task bar of the Administrative console window.
8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

## **ActivitySession service settings**

Use this page to administer the run-time properties of the ActivitySession service.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Activity Session Service**.

### **Startup**

Specifies whether the server will attempt to start the ActivitySession service.

Default                      Selected

Range

**Selected**

When the application server starts, it attempts to start the ActivitySession service automatically.

**Cleared**

The server does not try to start the ActivitySession service. If ActivitySessions are to be used in applications that run on this server, the system administrator must start the service manually or select this property then restart the server.

**Default timeout**

The default timeout for an ActivitySession. A server resets an ActivitySession if a remote client has failed to complete the ActivitySession within this time period.

The Default ActivitySession timeout specifies the time after which an ActivitySession is completed automatically by the ActivitySession service, if a remote client has failed to complete the ActivitySession within the specified time. The initial default timeout can be configured separately for each application server, and can be overridden programmatically by the UserActivitySession interface (setSessionTimeout).

|           |                               |
|-----------|-------------------------------|
| Data type | Integer                       |
| Units     | Seconds                       |
| Default   | 300 (5 minutes)               |
| Range     | -1 through 2147483647 seconds |

- -1 indicates that ActivitySessions never timeout
- 0 indicates that the default timeout applies
- Other values are an integer number of seconds

---

## Configuring the default ActivitySession timeout for an application server

Use this task to configure the default ActivitySession timeout for an application server, after which any started ActivitySessions are completed automatically by the ActivitySession service.

The ActivitySession timeout is used to reset any ActivitySession whose remote client has failed to complete the ActivitySession in a timely fashion. The initial default timeout can be configured separately for each application server, and can be overridden programmatically by the UserActivitySession interface (setSessionTimeout). If an ActivitySession that contains a transaction reaches the timeout, the transaction's timeout is accelerated so that it is timed out (and rolled back) immediately before the ActivitySession is reset.

To configure the default ActivitySession timeout for an application server, use the WebSphere Administrative console to complete the following steps:

Steps for this task

1. Start the WebSphere Administrative console.
2. In the navigation pane, expand **Servers-> Manage Application Servers**  
This displays a list of the application servers in the content pane.
3. In the Content pane, select the application server that you want to configure.

- This displays the properties for the application server in the content pane.
4. In the Additional Properties table, select **ActivitySession service**. This displays the ActivitySession service properties in the content pane.
  5. Set the **ActivitySession timeout** property to the default timeout as an integer number of seconds.
    - -1 indicates that ActivitySessions never timeout
    - 0 indicates that the default timeout, 300 seconds, applies
    - Other values are an integer number of seconds
  6. Click **OK**.
  7. To save your configuration, click **Save** on the task bar of the Administrative console window.
  8. **(Optional)** To have the changed configuration take effect, stop then restart the application server.

## ActivitySession service settings

Use this page to administer the run-time properties of the ActivitySession service.

To view this administrative console page, click **Servers > Application Servers > *server\_name* > Activity Session Service**.

### Startup

Specifies whether the server will attempt to start the ActivitySession service.

|         |  |
|---------|--|
| Default | Selected   |
| Range   | <p><b>Selected</b></p> <p>When the application server starts, it attempts to start the ActivitySession service automatically.</p> <p><b>Cleared</b></p> <p>The server does not try to start the ActivitySession service. If ActivitySessions are to be used in applications that run on this server, the system administrator must start the service manually or select this property then restart the server.</p> |

### Default timeout

The default timeout for an ActivitySession. A server resets an ActivitySession if a remote client has failed to complete the ActivitySession within this time period.

The Default ActivitySession timeout specifies the time after which an ActivitySession is completed automatically by the ActivitySession service, if a remote client has failed to complete the ActivitySession within the specified time. The initial default timeout can be configured separately for each application server, and can be overridden programmatically by the UserActivitySession interface (setSessionTimeout).

|           |   |
|-----------|---|
| Data type | Integer   |
| Units     | Seconds   |
| Default   | 300 (5 minutes)   |
| Range     | -1 through 2147483647 seconds <ul style="list-style-type: none"> <li>• -1 indicates that ActivitySessions never timeout</li> <li>• 0 indicates that the default timeout applies</li> <li>• Other values are an integer number of seconds</li> </ul> |

---

## Troubleshooting ActivitySessions

Use this overview task to help resolve a problem that you think is related to the ActivitySession service.

To identify and resolve ActivitySession-related problems, you can use the standard WebSphere Application Server RAS facilities. If you encounter a problem that you think might be related to ActivitySessions, complete the following stages:

### Steps for this task

1. Check for ActivitySession messages in the admin console.  
The ActivitySession service produces diagnostic messages prefixed by "WACS". The error message indicates the nature of the problem and provides some detail. The associated message information provides an explanation and any user actions to resolve the problem.
2. Check for ActivitySession messages in the activity log.  
Activity log messages produced by the ActivitySession service are accompanied by log analyzer descriptions.
3. Check for more messages in the application server's stdout.log.  
For more information about a problem, check the stdout.log file for the application server, which should contain more error messages and extra details about the problem.

---

## The ActivitySession service application programming interfaces

The ActivitySession service consists of an application programming interface available to Web applications, session EJBs, and J2EE client applications for application-managed demarcation of ActivitySession context. Applications use the UserActivitySession interface, which provides demarcation scope methods.

### ActivitySession API

The ActivitySession service provides the UserActivitySession interface for use by EJB Session beans using bean-managed context demarcation, Web application components configured with **ActivitySession control=Web Application**, and J2EE client applications. This UserActivitySession interface defines the set of ActivitySession operations available to an application component. An implementation of this interface is obtained via a JNDI lookup of the URL "java:comp/websphere/UserActivitySession". It is used to begin and end ActivitySessions and to query various attributes of the active ActivitySession associated with the thread.

For more information about the ActivitySession API, see the ActivitySession API (Javadoc).

The ActivitySession API and the implementation of its interfaces is contained in the com.ibm.websphere.ActivitySession package.

### Programming Examples

The following code extract provides a basic example of using the UserActivitySession interface:



```

// Get initial context
InitialContext ic = new InitialContext();
// Lookup UserActivitySession
UserActivitySession uas =
    (UserActivitySession)ic.lookup("java:comp/websphere/UserActivitySession");

// Set the ActivitySession timeout to 60 seconds
uas.setSessionTimeout(60);
// Start a new ActivitySession context
uas.beginSession();
// Do some work under this context
MyBeanA beanA.doSomething();
...
MyBeanB beanB.doSomethingElse();
// End the context
uas.endSession(EndModeCheckpoint);

```

---

## Samples: ActivitySessions

The following ActivitySession samples are provided with WebSphere Application Server:

### MasterMind sample

This sample is based on the game MasterMind. It consists of the following components:

- A servlet, configured with Activity session control set to Container, that accesses a stateful session bean.
- A stateful session bean, configured with an activation policy of ActivitySession containing transient state data.

The servlet begins an HttpSession at the start of each new game, and ends it at the end of each game; therefore an ActivitySession lasts for the duration of each game. The ActivitySession activation policy stops the bean from being passivated and therefore the transient data remains in memory. This is to demonstrate HttpSession/ActivationSession association in the web container, and an ActivitySession-scoped activation policy.

### J2EE client container application and a CMP entity bean backed by a 1-phase-commit datasource

In this sample, the entity bean is configured with the following properties:

- TX\_NOT\_SUPPORTED
- An ActivitySession container managed policy of REQUIRES
- An LTC boundary of ActivitySession
- An LTC Resolution Control of ContainerAtBoundary

The client accesses the UserActivitySession, begins an ActivitySession, updates two instances of the bean, then ends the ActivitySession. It does this twice using EndModeReset then EndModeCheckpoint. This sample demonstrates the following functionality:

- Client access to the UserActivitySession interface
- Multiple RMLTs being scoped to the ActivitySession and automatically taking their completion direction from that of the ActivitySession

The entity bean also holds a transient variable incremented by each method call (gets and sets for the persistent data). This value is checked before the end of the ActivitySession to show that the same bean instance is used. The client checks for the correct results.

## A J2EE client container application and two session beans with different ActivitySession types

This sample consists of a J2EE client container application and the following session beans:

- SLB1, a stateless session bean configured with an ActivitySession Type of Bean.
- SFB2, a stateful session bean configured with ActivitySession Type of Requires, an LTC boundary of ActivitySession, LTC Resolution Control of APPLICATION, and an LTC Unresolved Action of ROLLBACK.

Both beans are configured with TX\_NOTSUPPORTED.

This sample performs the following steps:

1. The client starts SLB1
2. SLB1 accesses the UserActivitySession interface, begins an ActivitySession, then calls a method on SFB2
3. SFB2 accesses the UserActivitySession interface, begins an ActivitySession, calls a method on SFB2
4. SFB2 gets a connection (setAutoCommit false) then uses JDBC to update a single-phase datasource.
5. SLB1 then optionally calls a separate method on SFB2 to finish the work either committing or rolling-back the RMLT.
6. SLB1 then ends the ActivitySession with an EndModeCheckpoint.

This demonstrates that the ActivitySession completion direction is unconnected to the direction of the RMLTs, although their containment is bound to the ActivitySession, and the use of the container using the unresolved action when the RMLT is not completed. It also shows a bean-managed ActivitySessions bean using the UserActivitySession interface. The sample checks for correct results and reports them back to the client.

---

## ActivitySession service: Resources for learning

Use the following links to find relevant supplemental information about ActivitySessions. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.



View links to additional information about:

- Programming model and decisions
- Programming specifications
- Other






### Programming model and decisions

-  ActivitySession API (Javadoc)

### Programming specifications

-  J2EE Activity Service for Extended Transactions
-  Java Transaction API (JTA) 1.0.1

### Other

-  WebSphere Application Server Enterprise Version 5 Overview: Advanced Transactional Connectivity
-  Listing of PDF files to learn about WebSphere Application Server Version 5
-  Listing of all IBM WebSphere Application Server Redbooks
-  Listing of all IBM WebSphere Application Server Whitepapers
-  WebSphere Application Server Enterprise Edition 4.0: A Programmer's Guide



---

## Chapter 24. Using naming

Naming is used by clients of WebSphere Application Server applications most commonly to obtain references to objects related to those applications, such as Enterprise JavaBeans (EJB) homes. The following steps outline the context of Naming in the overall application development and deployment process. Steps for this task follow:

### Steps for this task

1. Develop your application using either JNDI or CosNaming (CORBA) interfaces.  
Use these interfaces to look up server application objects that are bound into the name space and obtain references to them. Most Java developers use the JNDI interface. However, the CORBA CosNaming interface is also available for performing Naming operations on WebSphere Application Server name servers or other CosNaming name servers.
2. Assemble your application  
Application assembly is a packaging and configuration step that is a prerequisite to application deployment. If the application you are assembling is a client to an application running in another process, you should qualify the `jndiName` values in the deployment descriptors for the objects related to the other application. Otherwise, you may need to override the names with qualified names during application deployment. If the objects have fixed qualified names configured for them, you should use them so that the `jndiName` values do not depend on the other application's location within the topology of the cell.
3. Deploy your application.  
Put your assembled application onto the application server. If the application you are assembling is a client to an application running in another server process, be sure to qualify the `jndiName` values for the other application's server objects if they are not already qualified.  
For more information on qualified names, see [Lookup names support in deployment descriptors and thin clients](#).
4. Configure name space bindings.  
This step is necessary in these cases:
  - Your deployed application is to be accessed by legacy client applications running on previous versions of WebSphere Application Server. In this case, you must configure additional name bindings for application objects relative to the default initial context for legacy clients. (Version 5 clients have a different initial context from legacy clients.)
  - The application requires qualified name bindings for such reasons as:
    - It will be accessed by J2EE client applications or server applications running in another server process
    - It will be accessed by thin client applications

In this case, you can configure name bindings as additional bindings for application objects. The qualified names for the configured bindings are *fixed*, meaning they do not contain elements of the cell topology that can change if the application is moved to another server. Objects as bound into the name

space by the system can always be qualified with a topology-based name. You must explicitly configure a name binding to use as a fixed qualified name.

For more information on qualified names, see Lookup names support in deployment descriptors and thin clients. For more information on configured name bindings, see Configured name bindings.

5. Troubleshoot any problems that develop.

If a Naming operation is failing and you need to verify whether certain name bindings exist, use the `dumpNameSpace` tool to generate a dump of the name space.

---

## Naming

Naming is used by clients of WebSphere Application Server applications to obtain references to objects related to those applications, such as Enterprise JavaBeans (EJB) homes.

These objects are bound into a mostly hierarchical structure, referred to as a *name space*. In this structure, all non-leaf objects are called *contexts*. Leaf objects can be contexts and other types of objects. Naming operations, such as lookups and binds, are performed on contexts. All naming operations begin with obtaining an *initial context*. You can view the initial context as a starting point in the name space.

The name space structure consists of a set of *name bindings*, each consisting of a name relative to a specific context and the object bound with that name. For example, the name `myApp/myEJB` consists of one non-leaf binding with the name `myApp`, which is a context. The name also includes one leaf binding with the name `myEJB`, relative to `myApp`. The object bound with the name `myEJB` in this example happens to be an EJB home reference. The whole name `myApp/myEJB` is relative to the initial context, which you can view as a starting place when performing naming operations.

You can access and manipulate the name space through a *name server*. Users of a name server are referred to as *naming clients*. Naming clients typically use the Java Naming and Directory Interface (JNDI) to perform naming operations. Naming clients can also use the Common Object Request Broker Architecture (CORBA) `CosNaming` interface.

Typically, objects bound to the name space are resources and objects associated with installed applications. These objects are bound by the system, and client applications perform lookup operations to obtain references to them. Occasionally, server and client applications bind objects to the name space. An application can bind objects to transient or persistent partitions, depending on requirements.

In J2EE environments, some JNDI operations are performed with `java:` URL names. Names bound under these names are bound to a completely different name space which is local to the calling process. However, some lookups on the `java:` name space may trigger indirect lookups to the name server.

---

## New features for name space support

The following are new features of the WebSphere Application Server V5 naming implementation:

- **Name space is distributed.**

For additional scalability, the name space for a cell is distributed among various servers. Every server has a name server. In previous releases, there was only one name server for an entire administrative domain.

In WebSphere Application Server versions prior to V5, all servers shared the same default initial context, and everything was bound relative to that same initial context. In WebSphere Application Server V5, the default initial context for a server is its server root. System artifacts, such as EJB homes and resources, are bound to the server root of the server with which they are associated.

- **Transient and persistent partitions.**

The name space is partitioned into transient areas and persistent areas. Server roots are transient. System-bound artifacts such as EJB homes and resources are bound under server roots. There is a cell persistent root, which you can use for cell-scoped persistent bindings, and a node persistent root, which you can use to bind objects with a node scope.

- **System name space structure.**

The name space for the entire cell is federated among all servers in the cell. Every server process contains a name server. All name servers provide the same logical view of the cell name space. The various server roots and persistent partitions of the name space are interconnected by means of a system name space. You can use the system name space structure to traverse to any context in the cell name space.

- **Configured bindings.**

You can use the configuration graphical interface and script interfaces to configure bindings in various root contexts within the name space. These bindings are read-only and are bound by the system at server startup.

- **Support for CORBA Interoperable Naming Service (INS) object URLs.**

WebSphere Application Server V5 contains support for Common Object Request Broker Architecture (CORBA) object URLs (corbaloc and corbname) as Java Naming and Directory Interface (JNDI) provider URLs and lookup names.

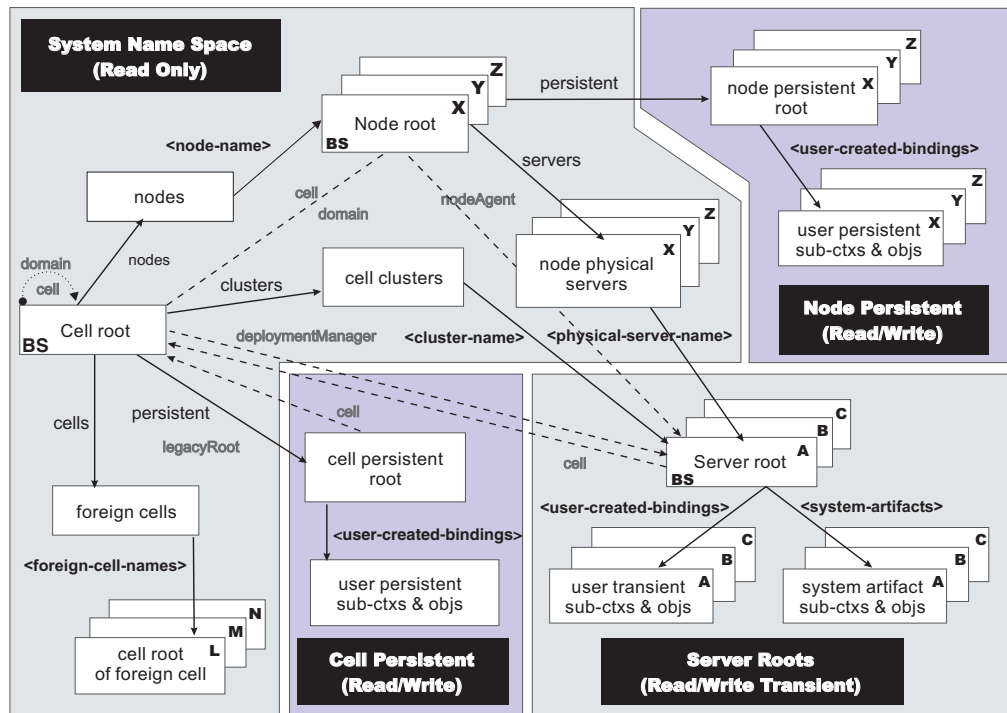
---

## Name space logical view

The name space for the entire cell is federated among all servers in the cell. Every server process contains a name server. All name servers provide the same logical view of the cell name space. The various server roots and persistent partitions of the name space are interconnected by a system name space. You can use the system name space structure to traverse to any context in a the cell's name space.

A logical view of the name space is shown in the following diagram.

### Logical View of a Cell's Name Space



### Name Space Logical View

The bindings in the preceding diagram appear with solid arrows, labeled in bold, and dashed arrows, labeled in gray. Solid arrows represent *primary bindings*. A primary binding is formed when the associated subcontext is created. Dashed arrows show *linked bindings*. A linked binding is formed when an existing context is bound under an additional name. Linked bindings are added for convenience or interoperability with previous WebSphere Application Server versions.

A cell name space is composed of contexts which reside in servers throughout the cell. All name servers in the cell provide the same logical view of the cell name space. A name server constructs this view at startup by reading configuration information. Each name server has its own local in-memory copy of the name space and does not require another running server to function. There are, however, a few exceptions. Server roots for other servers are not replicated among all the servers. The respective server for a server root must be running to access that server root context.

## Name space partitions

There are four major partitions in a cell name space:

- System name space partition
- Server roots partition
- Cell persistent partition
- Node persistent partition

Each partition is discussed in more detail below.

### System name space partition



The system name space contains a structure of contexts based on the cell topology. The system structure supports traversal to all parts of a cell name space and to the cell root of other cells, which are configured as foreign cells. The root of this structure is the cell root. In addition to the cell root, the system structure contains a node root for each node in the cell. You can access other contexts of interest specific to a node from the node root, such as the node persistent root and server roots for servers configured in that node.

All contexts in the system name space are read-only. You cannot add, update, or remove any bindings.

### **Server roots partition**

Each server in a cell has a server root context. A server root is specific to a particular server. You can view the server roots for all servers in a cell as being in a transient read/write partition of the cell name space. System artifacts, such as EJB homes for server applications and resources, are bound under the server root context of the associated server. A server application can also add bindings under its server root. These bindings are transient. Therefore, the server application creates all required bindings at application startup, so they exist anytime the application is running.

A server cluster is composed of many servers that are logically equivalent. Each member of the cluster has its own server root. These server roots are not replicated across the cluster. In other words, adding a binding to the server root of one member does not propagate it to the server roots of the other cluster members. To maintain the same view across the cluster, you should create all user bindings under the server root by the server application at application startup so that the bindings are present under the server root of each cluster member. Because of Workload Management (WLM) behavior, a JNDI client outside a cluster has no control over which cluster member's server root context becomes the target of the JNDI operation. Therefore, you should execute bind operations to the server root of a cluster member should from within that cluster member process only.

Distributing application objects among many server roots is a departure from previous WebSphere Application Server releases, where all system artifacts were bound under a single root. This change can affect the names that clients use to look up these objects.

Server-scoped bindings are relative to a server's server root.

### **Cell persistent partition**

The root context of the cell persistent partition is the cell persistent root. A binding created under the cell persistent root is saved as part of the cell configuration and continues to exist until it is explicitly removed. Applications that need to create additional persistent bindings of objects generally associated with the cell can bind these objects under the cell persistent root.

It is important to note that the cell persistent area is not designed for transient, rapidly changing bindings. The bindings are more static in nature, such as part of an application setup or configuration, and are not created at run time.

An important role of the cell persistent root is as the initial context for clients running in previous WebSphere Application Server versions. If you want to access an enterprise bean by WebSphere Application Server v4.0.x and 3.5.x clients, you

must ensure that a binding for it has been added to the cell persistent root. You can configure these additional bindings as cell-scoped bindings.

### Node persistent partition

The node persistent partition is similar to the cell partition except that each node has its own node persistent root. A binding created under a node persistent root is saved as part of that node configuration and continues to exist until it is explicitly removed.

Applications that need to create additional persistent bindings of objects associated with a specific node can bind those objects under that particular node's node persistent root. As with the cell persistent area, it is important to note that the node persistent area is not designed for transient, rapidly changing bindings. These bindings are more static in nature, such as part of an application setup or configuration, and are not created at run time.

Unlike the cell persistent root, the node persistent root plays no special role in interoperability with WebSphere Application Server clients of previous releases. Node-scoped bindings are relative to a node's node persistent root.

---

## Initial context support

All naming operations begin with obtaining an initial context. You can view the initial context as a starting point in the name space. Use the initial context to perform naming operations, such as looking up and binding objects in the name space.

### Initial contexts registered with the ORB as initial references

The server root, cell persistent root, cell root, and node root are registered with the name server's ORB and can be used as an initial context. An initial context is used by CORBA and enterprise bean applications as a starting point for name space lookups. The keys for these roots as recognized by the ORB are shown in the following table:

| Root Context         | Initial Reference Key              |
|----------------------|------------------------------------|
| Server Root          | NameServiceServerRoot              |
| Cell Persistent Root | NameServiceCellPersistentRoot      |
| Cell Root            | NameServiceCellRoot<br>NameService |
| Node Root            | NameServiceNodeRoot                |

A server root initial context is the server root context for the specific server you are accessing. Similarly, a node root initial context is the node root for the server being accessed.

You can use the previously mentioned keys in CORBA INS object URLs (corbaloc and corbaname) and as an argument to an ORB `resolve_initial_references` call. For examples, see CORBA and JNDI programming examples, which show how to get an initial context.

## Default initial contexts

The default initial context depends on the type of client. Different categories of clients and the corresponding default initial context follow.

### WebSphere Application Server V5 JNDI interface implementation

The JNDI interface is used by EJB applications to perform name space lookups. WebSphere Application Server clients by default use the WebSphere Application Server CosNaming JNDI plug-in implementation. The default initial context for clients of this type is the server root of the server specified by the provider URL. For more details, refer to the JNDI programming examples on getting initial contexts.

### WebSphere Application Server JNDI interface implementation prior to V5

WebSphere Application Server clients running in releases prior to WebSphere Application Server V5 by default use WebSphere Application Server's v4.0 CosNaming JNDI plug-in implementation. The default initial context for clients of this type is the cell persistent root, also known as the *legacy root*.

### Other JNDI implementation

Some applications can perform name space lookups with a non-WebSphere Application Server CosNaming JNDI plug-in implementation. Assuming the key **NamingContext** is used to obtain the initial context, the default initial context for clients of this type is the cell root.

### CORBA

The standard CORBA client obtains an initial `org.omg.CosNaming.NamingContext` reference with the key **NamingContext**. The initial context in this case is the cell root.

---

## Lookup names support in deployment descriptors and thin clients

Server objects, such as EJB homes, are bound relative to the server root context for the server in which the application is installed. Other objects, such as resources, can also be bound to a specific server root. The names used to look up these objects must be qualified so as to select the correct server root. This is a departure from previous versions of WebSphere Application Server, where these objects were all bound under a single root context. This section discusses what relative and qualified names are, when they can be used, and how you can construct them.

### Relative names

All names are relative to a context. Therefore, a name that can be resolved from one context in the name space cannot necessarily be resolved from another context in the name space. This point is significant because the system binds objects with names relative to the server root context of the server in which the application is installed. Each server has its own server root context. The initial JNDI context is by default the server root context for the server identified by the provider URL used to obtain the initial context. (Typically, the URL consists of a host and port.) For applications running in a server process, the default initial JNDI context is the server root for that server. A relative name will resolve successfully when the initial context is obtained from the server which contains the target object, but it will not resolve successfully from an initial context obtained from another server.

If all clients of a server application run in the same server process as the application, all objects associated with that application are bound to the same initial context as the clients' initial context. In this case, only names relative to the server's server root context are required to access these server objects. Frequently, however, a server application has clients that run outside the application's server process. The initial context for these clients can be different from the server application's initial context, and lookups on the relative names for server objects may fail. These clients need to use the qualified name for the server objects. This point must be considered when setting up the `jndiName` values in a J2EE client application deployment descriptors and when constructing lookup names in thin clients. Qualified names resolve successfully from any initial context in the cell.

## Qualified names

All names are relative to a context. Here, the term qualified name refers to names that can be resolved from any initial context in a cell. This action is accomplished by using names that navigate to the same context, the cell root. The rest of the qualified name is then relative to the cell root and uniquely identifies an object throughout the cell. All initial contexts in a server (that is, all naming contexts in a server registered with the ORB as an initial reference) contain a binding with the name `cell`, which links back to the cell root context. All qualified names begin with the string `cell/` to navigate from the current initial context back to the cell root context.

A qualified name for an object is the same throughout the cell. The name can be topology-based, or some fixed name bound under the cell persistent root. Topology-based names, described in more detail below, navigate through the system name space to reach the target object. A fixed name bound under the cell persistent root has the same qualified name throughout the cell and is independent of the topology. Creating a fixed name under the cell persistent root for a server application object requires an extra step when the server application is installed, but this step eliminates impacts to clients when the application is moved to a different location in the cell topology. The process for creating a fixed name is described later in this section.

Generally speaking, you **must** use qualified names for EJB `jndiName` values in a J2EE client application deployment descriptors and for EJB lookup names in thin clients. The only exception is when the initial context is obtained from the server in which the target object resides. For example, a session bean which is a client to an entity bean can use a relative name if the two beans run in the same server. If the session bean and entity beans run in different servers, the `jndiName` for the entity bean must be qualified in the session bean's deployment descriptors. The same requirement may be true for resources as well, depending on the scope of the resource.

### Topology-based names

The system name space partition in a cell's name space reflects the cell's topology. This structure can be navigated to reach any object bound into the cell's name space. Topology-based qualified names include elements from the topology which reflect the object's location within the cell. For a system-bound object, such as an EJB home, the form for a topology-based qualified name depends on whether the object is bound to a single server or cluster. Both forms are described below.

#### Single Server

An object bound in a single server has a topology-based qualified name of the following form:

`cell/nodes/nodeName/servers/serverName/relativeJndiName`

where *nodeName* and *serverName* are the node name and server name for the server where the object is bound, and *relativeJndiName* is the unqualified name of the object; that is, the object's name relative to its server's server root context.

### Server Cluster

An object bound in a server cluster has a topology-based qualified name of the following form:

`cell/clusters/clusterName/relativeJndiName`

where *clusterName* is the name of the server cluster where the object is bound, and *relativeJndiName* is the unqualified name of the object; that is, the object's name relative to a cluster member's server root context.

### Fixed names

It is possible to create a fixed name for a server object so that the qualified name is independent of the cell topology. This quality is desirable when clients of the application run in other server processes or as pure clients. Fixed names have the advantage of not changing if the object is moved to another server. The `jndiName` values in deployment descriptors for a J2EE client application can reference the qualified fixed name for a server object regardless of the cell topology on which the client or server application is being installed.

Defining a cell-wide fixed name for a server application object requires an extra step after the server application is installed. That is, a binding for the object must be created under the cell persistent root. A fixed name bound under the cell persistent root can be any name, but all names under the cell persistent root must be unique within the cell because the cell persistent root is global to the entire cell.

A qualified fixed name has the form:

`cell/persistent/fixedName`

where *fixedName* is an arbitrary fixed name.

The binding can be created programmatically (for example, using JNDI). However, it is probably more convenient to configure a cell-scoped binding for the server object.

You must keep the programmatic or configured binding up-to-date. Configured EJB bindings are based on the location of the enterprise bean within the cell topology, and moving the EJB application to another single server or to a server cluster, for example, requires the configured binding to be updated. Similar changes affect an EJB home reference programmatically bound so that the fixed name would need to be rebound with a current reference. However, for J2EE clients, the `jndiName` value for the object, and for thin clients, the lookup name for the object, remains the same. In other words, clients that access objects by fixed names are not affected by changes to the configuration of server applications they access.

---

## JNDI support in WebSphere Application Server

IBM WebSphere Application Server includes a name server to provide shared access to Java components, and an implementation of the `javax.naming` JNDI package which supports user access to the WebSphere Application Server name server through the JNDI naming interface.

WebSphere Application Server does *not* provide implementations for:

- `javax.naming.directory` or
- `javax.naming.ldap` packages

Also, WebSphere Application Server does *not* support interfaces defined in the `javax.naming.event` package.

However, to provide access to LDAP servers, the development kit shipped with WebSphere Application Server supports Sun's implementation of:

- `javax.naming.ldap` and
- `com.sun.jndi.ldap.LdapCtxFactory`

WebSphere Application Server's JNDI implementation is based on version 1.2 of the JNDI interface, and was tested with Version 1.2.1 of Sun's JNDI Service Provider Interface (SPI).

The default behavior of this JNDI implementation is adequate for most users. However, users with specific requirements can control certain aspects of JNDI behavior.

---

## Developing applications that use JNDI

References to EJB homes and other artifacts such as data sources are bound to the WebSphere name space. These objects can be obtained through the JNDI interface. Before you can perform any JNDI operations, you need to get an initial context. You can use the initial context to look up objects bound to the WebSphere name space.

These examples describe how to get an initial context and how to perform lookup operations.

- Getting the default initial context
- Getting an initial context by setting the provider URL property
- Setting the provider URL property to select a different root context as the initial context
- Looking up an EJB home with JNDI
- Looking up a JavaMail session with JNDI

In these examples, the default behavior of features specific to WebSphere's JNDI Context implementation is used.

WebSphere Application Server's JNDI context implementation includes special features. JNDI caching enhances performance of repeated lookup operations on the same objects. Name syntax options offer a choice of a name syntaxes, one optimized for typical JNDI clients, and one optimized for interoperability with CosNaming applications. Most of the time, the default behavior of these features is the preferred behavior. However, sometimes you should modify the behavior for specific situations.

JNDI caching and name syntax options are associated with a `javax.naming.InitialContext` instance. To select options for these features, set properties that are recognized by the WebSphere Application Server's initial context factory. To set JNDI caching or name syntax properties which will be visible to WebSphere Application Server's initial context factory, follow the following steps.

#### Steps for this task

##### 1. **(Optional)** Configure JNDI caches

JNDI caching can greatly increase performance of JNDI lookup operations. By default, JNDI caching is enabled. In most situations, this default is the desired behavior. However, in specific situations, use the other JNDI cache options.

Objects are cached locally as they are looked up. Subsequent lookups on cached objects are resolved locally. However, cache contents can become stale. This situation is not usually a problem, since most objects you look up do not change frequently. If you need to look up objects which change relatively frequently, change your JNDI cache options.

JNDI clients can use several properties to control cache behavior.

You can set properties:

- From the command line by entering the actual string value. For example:  
`java -Dcom.ibm.websphere.naming.jndicache.maxentrylife=1440`
- In a `jndi.properties` file by creating a file named `jndi.properties` as a text file with the desired properties settings. For example:

```
...
com.ibm.websphere.naming.jndicache.cacheobject=none
...
```

Include the file as the beginning of the classpath, so that the classloader loads your copy of `jndi.properties` before any other copies.

- Within a Java program by using the **PROPS.JNDI\_CACHE\*** Java constants, defined in the `com.ibm.websphere.naming.PROPS` file. The constant definitions follow:

```
public static final String JNDI_CACHE_OBJECT = "com.ibm.websphere.naming.jndicache.cacheobject";
public static final String JNDI_CACHE_OBJECT_NONE = "none";
public static final String JNDI_CACHE_OBJECT_POPULATED = "populated";
public static final String JNDI_CACHE_OBJECT_CLEARED = "cleared";
public static final String JNDI_CACHE_OBJECT_DEFAULT = JNDI_CACHE_OBJECT_POPULATED;

public static final String JNDI_CACHE_NAME = "com.ibm.websphere.naming.jndicache.cachename";
public static final String JNDI_CACHE_NAME_DEFAULT = "providerURL";

public static final String JNDI_CACHE_MAX_LIFE = "com.ibm.websphere.naming.jndicache.maxcachelife";
public static final int JNDI_CACHE_MAX_LIFE_DEFAULT = 0;

public static final String JNDI_CACHE_MAX_ENTRY_LIFE = "com.ibm.websphere.naming.jndicache.maxentrylife";
public static final int JNDI_CACHE_MAX_ENTRY_LIFE_DEFAULT = 0;
```

To use the previous properties in a Java program, add the property setting to a hashtable and pass it to the `InitialContext` constructor as follows:

```
java.util.Hashtable env = new java.util.Hashtable();
...
//Disable caching
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_NONE);
...
javax.naming.Context initialContext = new javax.naming.InitialContext(env);
```

##### 2. **(Optional)** Specify the name syntax

Most WebSphere applications use JNDI to look up EJB objects and do not need to look up objects bound by CORBA applications. Therefore, the default name syntax used for JNDI names is the most convenient. If your application needs to look up objects bound by CORBA applications, you may need to change your name syntax so that all CORBA CosNaming names can be represented.

JNDI clients can set the name syntax by setting a property. The property setting is applied by the initial context factory when you instantiate a new `java.naming.InitialContext` object. Names specified in JNDI operations on the initial context are parsed according to the specified name syntax.

You can set the property:

- From the command line by entering the actual string value. For example:

```
java -Dcom.ibm.websphere.naming.name.syntax=ins
```

- In a `jndi.properties` file by creating a file named `jndi.properties` as a text file with the desired properties settings. For example:

```
...
com.ibm.websphere.naming.name.syntax=ins
...
```

Include the file as the beginning of the classpath, so that the classloader loads your copy of `jndi.properties` before any other copies.

- Within a Java program by using the `PROPS.NAME_SYNTAX*` Java constants, defined in the `com.ibm.websphere.naming.PROPS` file. The constant definitions follow:

```
public static final String NAME_SYNTAX="com.ibm.websphere.naming.name.syntax";
public static final String NAME_SYNTAX_JNDI = "jndi";
public static final String NAME_SYNTAX_INS = "ins";
```

To use the previous properties in a Java program, add the property setting to a hashtable and pass it to the `InitialContext` constructor as follows:

```
java.util.Hashtable env = new java.util.Hashtable();
...
env.put(PROPS.NAME_SYNTAX, PROPS.NAME_SYNTAX_INS); // Set name syntax to INS
...
javax.naming.Context initialContext = new javax.naming.InitialContext(env);
```

## Example: Getting the default initial context

This example below gets the default initial context. That is, no provider URL is passed to the `javax.naming.InitialContext` constructor. The following section explains the process of determining the address of the bootstrap server to use to obtain the initial context.

### Usage scenario

```
...
import javax.naming.Context;
import javax.naming.InitialContext;
...
Context initialContext = new InitialContext();
...
```

The default initial context returned depends the runtime environment of the JNDI client. The initial context returned in the various environments are listed below:

- Thin client: The server root context of the server running on the local host at port 2809.
- Pure client:



- The context specified by the `java.naming.provider.url` property passed to `launchClient` command with the `-CCD` command line parameter. The context usually will be the server root context of the server at the address specified in the URL, although it is possible to construct a `corbaname` or `corbaloc` URL which resolves to some other context.
- If no provider URL was specified, the server root context of the server running on the host and port specified by the `-CCBootstrapHost` `-CCBootstrapPort` command line parameters. The default host is the local host, and the default port is 2809.
- Server process: The server root context for that process.

Even though no provider URL is explicitly specified in the above example, the `InitialContext` may find a provider URL defined in other places that it searches for property settings.

Users of properties which affect ORB initialization should read the rest of this section for a deeper understanding of exactly how initial contexts are obtained, which has changed from previous releases.

### **Determining which server is used to obtain the initial context**

WebSphere Application Server name servers are CORBA CosNaming name servers, and WebSphere Application Server provides a CosNaming JNDI plug-in implementation for JNDI clients to perform naming operations on WebSphere Application Server name spaces. The WebSphere Application Server CosNaming plug-in implementation is selected through a JNDI property that is passed to the `InitialContext` constructor. This property is `java.naming.factory.initial`, and it specifies the initial context factory implementation to use to obtain an initial context. The factory returns a `javax.naming.Context` instance, which is part of its implementation.

The WebSphere Application Server initial context factory, `com.ibm.websphere.naming.WsnInitialContextFactory`, is typically used by WebSphere Application Server applications to perform JNDI operations. The WebSphere Application Server run-time environment is set up to use this WebSphere Application Server initial context factory if one is not specified explicitly by the JNDI client. When the initial context factory is invoked, an *initial context* is obtained. The following paragraphs explain how the WebSphere Application Server initial context factory obtains the initial context in client and server environments.

### **Understanding the registration of initial references in server processes**

Every WebSphere Application Server has an ORB used to receive and dispatch invocations on objects running in that server. Services running in the server process can register initial references with the ORB. Each initial reference is registered under a key, which is a string value. An initial reference can be any CORBA object. WebSphere Application Server name servers register several initial contexts as initial references under predefined keys. Each name server initial reference is an instance of the interface `org.omg.CosNaming.NamingContext`.

### **Obtaining initial references in pure client processes**

Pure JNDI clients, that is, JNDI clients which are not running in a WebSphere Application Server process, also have an ORB instance. This client ORB instance can be passed to the `InitialContext` constructor, but typically the initial context factory creates and initializes the client ORB instance transparently. A client ORB

can be initialized with initial references, but the initial references most likely resolve to objects running in some server. The initial context factory does not define any default initial references when it initializes an ORB. If the `resolve_initial_references` method is invoked on the client ORB when no initial references have been configured, the method invocation fails. This condition is typical for pure client processes. To obtain an initial NamingContext reference, the initial context factory must invoke `string_to_object` with an IIOP type CORBA object URL, such as `corbaloc:iiop:myhost:2809`. The URL specifies the address of the server from which to obtain the initial context. The host and port information is extracted from the provider URL passed to the InitialContext constructor. If no provider URL is defined, the WebSphere Application Server initial context factory uses the default provider URL of `corbaloc:iiop:localhost:2809`. The `string_to_object` ORB method resolves the URL and communicates with the target server ORB to obtain the initial reference.

### Obtaining initial references in server processes

If the JNDI client is running in a WebSphere Application Server process, the initial context factory obtains a reference to the server ORB instance if the JNDI client does not provide an ORB instance. Typically, JNDI clients running in server processes use the server ORB instance; that is, they do not pass an ORB instance to the InitialContext constructor. The name server which is running in the server process sets a provider URL as a `java.lang.System` property to serve as the default provider URL for all JNDI clients in the process. This default provider URL is `corbaloc:rir:/NameServiceServerRoot`. This URL resolves to the server root context for that server. (The URL is equivalent to invoking `resolve_initial_references` on the ORB with a key of `NameServiceServerRoot`. The name server registers the server root context as an initial reference under that key.)

### Understanding the legacy ORB protocol

Previous versions of WebSphere Application Server used a different ORB implementation, which used a legacy protocol in contrast with the Interoperable Name Service (INS) protocol now used. This change has affected the implementation of the WebSphere Application Server initial context factory. **Certain types of pure clients can experience different behavior when getting initial JNDI contexts as compared to previous releases of WebSphere Application Server.** This behavior is discussed in more detail below.

The following ORB properties are used with the legacy ORB protocol for ORB initialization and are now deprecated:

- `com.ibm.CORBA.BootstrapHost`
- `com.ibm.CORBA.BootstrapPort`

The new INS ORB is different in a major respect, in that it exhibits no default behavior if no initial references are defined. In the legacy ORB, the bootstrap host and port values defaulted to `localhost` and `900`. All initial references were obtained from the server running on the bootstrap host and port. So, if the ORB user provided no bootstrap host and port, all initial references are resolved from the server running on the local host at port `900`. The INS ORB has no concept of bootstrap host or bootstrap port. All initial references are defined independently. That is, different initial references could resolve to different servers. If `ORB.resolve_initial_references` is invoked with a key such that the ORB is not initialized with an initial reference having that key, the call fails.

In previous releases of WebSphere Application Server, the initial context factory invoked `resolve_initial_references` on the ORB in the absence of any provider URL. This action succeeded if a name server at the default bootstrap host and port was running. Today, with the INS ORB, this would fail. (Actually, the ORB would fall back to the legacy protocol during the deprecation period, but when the legacy protocol is no longer supported, the operation would fail.) The initial context factory now uses a default provider URL of `corbaloc:iiop:localhost:2809`, and invokes `string_to_object` with the provider URL. This operation preserves the behavior that pure clients in previous releases experienced when they set no ORB bootstrap properties or provider URL. **However, this different initial context factory implementation changes the behavior experienced by certain legacy pure clients, which do not specify a provider URL:**

- Clients which set the ORB bootstrap properties listed above when getting an initial context.
- Clients which supply their own ORB instance to the `InitialContext` constructor.

There are two ways to circumvent this change of behavior:

- Always specify an IIOP type provider URL. This approach does not depend on the bootstrap host and port properties and continues to work when support for the bootstrap host and port properties is removed. For example, you can express bootstrap host and port property values of `myHost` and `2809`, respectively, as `corbaloc:iiop:myHost:2809`.
- Use an `rir` type provider URL:
  - Specify `corbaloc:rir:/NameServiceServerRoot` if the ORB is initialized to use a WebSphere Application Server 5 server as the bootstrap server.
  - Specify `corbaname:rir:/NameService#domain/legacyRoot` if the ORB is initialized to use a WebSphere Application Server 4.0.x server as the bootstrap server.
  - Specify `corbaloc:rir:/NameService` if the ORB is initialized to use a server other than a WebSphere Application Server 5 or 4.0.x server as the bootstrap server.

URLs of this type are equivalent to invoking `resolve_initial_references` on the ORB with the specified key. If the bootstrap host and port properties are being used to initialize the ORB, this approach will not work when the bootstrap and host properties are no longer supported.

### The `InitialContext` constructor search order for JNDI properties

If the code snippet shown at the beginning of this section is executed by an application, the bootstrap server depends on the value of the property, `java.naming.provider.url`. If the property is not set (in server processes the default value is set as a system property), the default host of `localhost` and default port of `2809` are used as the address of the server from which to obtain the initial context. The JNDI specification describes where the `InitialContext` constructor looks for `java.naming.provider.url` property settings, but briefly, the property is picked up from the following places in the order shown:

1. The `InitialContext` constructor. This does not apply to the above example since the example uses the empty `InitialContext` constructor.
2. System environment. You can add JNDI properties to the system environment as an option on the `java` command invocation and by program code. The recommended way to set the provider URL in the system environment is as an option supplied to the `Java` command invocation. Setting the provider URL in this manner is not temporal, so that getting a default initial context will always

yield the same result. It is generally recommended that program code not set the provider URL property in the system environment because as a side-effect, this could adversely affect other, possibly unrelated, code running elsewhere in the same process.

3. `jndi.properties` file. There may be many `jndi.properties` files that are within the scope of the class loader in effect. All `jndi.properties` files are used for setting JNDI properties, but the provider URL setting is determined by the first `jndi.properties` file returned by the class loader.

## Example: Getting an initial context by setting the provider URL property

In general, JNDI clients should assume the correct environment is already configured so there is no need to explicitly set property values and pass them to the `InitialContext` constructor. However, a JNDI client may need to access a name space other than the one identified in its environment. In this case, it is necessary to explicitly set the `java.naming.provider.url` (provider URL) property used by the `InitialContext` constructor. A provider URL contains bootstrap server information that the initial context factory can use to obtain an initial context. Any property values passed in directly to the `InitialContext` constructor take precedence over settings of those same properties found elsewhere in the environment.

You can use two different provider URL forms with WebSphere Application Server's initial context factory:

- A CORBA object URL (new for J2EE 1.3)
- An IIOP URL

CORBA object URLs are more flexible than IIOP URLs and are the recommended URL format to use. CORBA object URLs are part of the OMG CosNaming Interoperable Naming Specification. A `corbaname` URL, for example, can include initial context and lookup name information and can be used as a lookup name without the need to explicitly obtain another initial context. The IIOP URLs are the legacy JNDI format, but are still supported by the WebSphere Application Server initial context factory.

The following examples illustrate the use of these URLs.

### Using a CORBA object URL

This example shows a CORBA object URL.

#### Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL,
        "corbaloc:iiop:myhost.mycompany.com:2809");
Context initialContext = new InitialContext(env);
...
```

### Using a CORBA object URL with multiple name server addresses

CORBA object URLs can contain more than one bootstrap address. You can use this feature when attempting to obtain an initial context from a server cluster. You can specify the bootstrap addresses for all servers in the cluster in the URL. The

operation succeeds if at least one of the servers is running, eliminating a single point of failure. There is no guarantee of any particular order in which the address list will be processed. For example, the second bootstrap address may be used to obtain the initial context even though the server at the first bootstrap address in the list is available.

Multiple-address provider URLs should only contain the bootstrap addresses of members of the same cluster. Otherwise, incorrect behavior may occur.

An example of a corbaloc URL with multiple addresses follows.

#### Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
// All of the servers in the provider URL below are members of the same cluster.
env.put(Context.PROVIDER_URL,
        "corbaloc::myhost1:9810,:myhost1:9811,:myhost2:9810");
Context initialContext = new InitialContext(env);
...
```

### **Using a CORBA object URL from an non-WebSphere Application Server JNDI implementation**

Initial context factories for CosNaming JNDI plug-in implementations other than the WebSphere Application Server initial context factory most likely obtain an initial context using the object key, `NameService`. When you use such a context factory to obtain an initial context from a WebSphere Application Server name server, the initial context is the cell root context. Since system artifacts such as EJB homes associated with a server are bound under the server's server root context, names used in JNDI operations must be qualified. If you want to use relative names, ensure your initial context is the server root context under which the target object is bound. In order to make the server root context the initial context, specify a corbaloc provider URL with an object key of `NameServiceServerRoot`.

This example shows a CORBA object type URL from a non-WebSphere Application Server JNDI implementation. This example assumes full CORBA object URL support by the non-WebSphere Application Server JNDI implementation. The object key of `NameServiceServerRoot` is specified so that the initial context will be the specified server's server root context.

#### Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.somecompany.naming.TheirInitialContextFactory");
env.put(Context.PROVIDER_URL,
        "corbaname:iiop:myhost.mycompany.com:9810/NameServiceServerRoot");
Context initialContext = new InitialContext(env);
...
```

If qualified names are used, you can use the default key of `NameService`.

## Using an IIOP URL

The IIOP type of URL is a legacy format which is not as flexible as CORBA object URLs. However, URLs of this type are still supported. The following example shows an IIOP type URL as the provider URL.

### Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, "iiop://myhost.mycompany.com:2809");
Context initialContext = new InitialContext(env);
...
```

## Example: Setting the provider URL property to select a different root context as the initial context

Each server contains its own server root context, and, when bootstrapping to a server, the server root is the default initial JNDI context. Most of the time, this default is the desired initial context, since system artifacts such as EJB homes are bound there. However, other root contexts exist, which can contain bindings of interest. It is possible to specify a provider URL to select other root contexts.

### Selecting the initial root context with a CORBA object URL

There are several object keys registered with the bootstrap server that you can use to select the root context for the initial context. To select a particular root context with a CORBA object URL object key, set the object key to the corresponding value. The default object key is NameService. Using JNDI yields the server root context. A table that lists the different root contexts and their corresponding object key follows:

| Root Context         | CORBA Object URL Object Key   |
|----------------------|-------------------------------|
| Server Root          | NameServiceServerRoot         |
| Cell Persistent Root | NameServiceCellPersistentRoot |
| Cell Root            | NameServiceCellRoot           |
| Node Root            | NameServiceNodeRoot           |

The following example shows the use of a corbaloc URL with the object key set to select the cell persistent root context as the initial context.

### Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL,
        "corbaloc:iiop:myhost.mycompany.com:2809/NameServiceCellPersistentRoot");
Context initialContext = new InitialContext(env);
...
```

## Selecting the initial root context with the name space root property

You can also select the initial root context by passing a name space root property setting to the InitialContext constructor. Generally, the object key setting described above is sufficient. Sometimes a property setting is preferable. For example, you can set the root context property on the Java invocation to make which server root is being used as the initial context transparent to the application. The default server root property setting is defaultroot, which yields the server root context.

| Root Context         | Name Space Root Property Value |
|----------------------|--------------------------------|
| Server Root          | bootstrapservroot              |
| Cell Persistent Root | cellpersistentroot             |
| Cell Root            | cellroot                       |
| Node Root            | bootstrapnoderoot              |

The initial context factory ignores the name space root property if the provider URL contains an object key other than NameService.

The following example shows use of the name space root property to select the cell persistent root context as the initial context. Note that available constants are used instead of hardcoding the property name and value.

### Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
import com.ibm.websphere.naming.PROPS;
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, "corbaloc:iiop:myhost.mycompany.com:2809");
env.put(PROPS.NAME_SPACE_ROOT, PROPS.NAME_SPACE_ROOT_CELL_PERSISTENT);
Context initialContext = new InitialContext(env);
...
```

## Example: Looking up an EJB home with JNDI

Most applications which use JNDI run in a container. Some do not. The name used to look up an object depends on whether or not the application is running in a container. The examples below show lookups from each type of application. Sometimes it is more convenient for an application to use a corbaname URL as the lookup name. Container-based JNDI clients and thin Java clients can use a corbaname URL. An example of a lookup with a corbaname URL is also included in this section.

### JNDI lookup from an application running in a container

Applications that run in a container can use java: lookup names. Lookup names of this form provide a level of indirection such that the lookup name used to look up an object is not dependent on the object's name as it is bound in the name server's name space. The deployment descriptors for the application provide the mapping from the java: name and the name server lookup name. The container sets up the java: name space based on the deployment descriptor information so that the java: name is correctly mapped to the corresponding object.

The following example shows a lookup of an EJB home. The actual home lookup name is determined by the application's deployment descriptors.

```

// Get the initial context as shown in a previous example
...
// Look up the home interface using the JNDI name
try {
    java.lang.Object ejbHome =
        initialContext.lookup("java:comp/env/com/mycompany/accounting/AccountEJB");
    accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
    ...
}

```

## JNDI lookup from an application that does not run in a container

Applications that do not run in a container cannot use `java:` lookup names because it is the container which sets the `java:` name space up for the application. Instead, an application of this type must look the object up directly from the name server. Each application server contains a name server. System artifacts such as EJB homes are bound relative to the server root context in that name server. The various name servers are federated by means of a system name space structure. The recommended way to look up objects on different servers is to qualify the name so that the name resolves from any initial context in the cell. If a relative name is used, the initial context must be the same server root context as the one under which the object is bound. The form of the qualified name depends on whether the qualified name is a topology-based name or a fixed name. A topology based name depends on whether the object resides in a single server or a server cluster. Examples of each form of qualified name follow.

### Topology-based qualified names

Topology-based qualified names traverse through the system name space to the server root context context under which the target object is bound. A topology-based qualified name resolves from any initial context in the cell. The topology-based qualified name depends on whether the object resides on a single server or server cluster. Examples of each lookup follow.

#### Single server

The following example shows a lookup of an EJB home that is running in the single server, `MyServer`, configured in the node, `Node1`.

```

// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the JNDI name
try {
    java.lang.Object ejbHome = initialContext.lookup(
        "cell/nodes/Node1/servers/MyServer/com/mycompany/accounting/AccountEJB");
    accountHome = (AccountHome)javax.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
    ...
}

```

#### Server cluster

The example below shows a lookup of an EJB home which is running in the cluster, `MyCluster`. The name can be resolved if any of the cluster members is running.

```

// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.

```



```

...
// Look up the home interface using the JNDI name
try {
    java.lang.Object ejbHome = initialContext.lookup(
        "cell/clusters/MyCluster/com/mycompany/accounting/AccountEJB");
    accountHome = (AccountHome)jvax.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
    ...
}

```

### Fixed qualified names

If the target object has a cell-scoped fixed name defined for it, you can use its qualified form instead of the topology-based qualified name. Even though the topology-based name works, the fixed name does not change with the specific cell topology or with the movement of the target object to a different server. An example lookup with a qualified fixed name is shown below.

```

// Get the initial context as shown in a previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the JNDI name
try {
    java.lang.Object ejbHome = initialContext.lookup(
        "cell/persistent/com/mycompany/accounting/AccountEJB");
    accountHome = (AccountHome)jvax.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
    ...
}

```

### JNDI lookup with a corbaname URL

A corbaname can be useful at times as a lookup name. If, for example, the target object is not a member of the federated name space and cannot be located with a qualified name, a corbaname can be a convenient way to look up the object. A lookup with a corbaname URL follows.

```

// Get the initial context as shown in a previous example
...
// Look up the home interface using a corbaname URL
try {
    java.lang.Object ejbHome =
        initialContext.lookup("corbaname:iiop:someHost:2809#com/mycompany/accounting/AccountEJB");
    accountHome = (AccountHome)jvax.rmi.PortableRemoteObject.narrow(
        (org.omg.CORBA.Object) ejbHome, AccountHome.class);
}
catch (NamingException e) { // Error getting the home interface
    ...
}

```

## Example: Looking up a JavaMail session with JNDI

The example below shows a lookup of a JavaMail resource. The actual lookup name is determined by the application's deployment descriptors.

### Usage scenario

```

// Get the initial context as shown above
...
Session session = (Session) initialContext.lookup("java:comp/env/mail/MailSession");

```

## JNDI interoperability considerations

This section explains considerations to take into account when interoperating with previous releases of WebSphere Application Server and with non-WebSphere Application Server JNDI clients. Also, the way resources from MQSeries must be bound to the name space has changed and is described below.

### Interoperability with previous WebSphere Application Server Releases

#### EJB clients running on WebSphere Application Server V3.5 or V4.0 accessing EJB applications running on WebSphere Application Server V5

Applications migrated from previous versions of WebSphere Application Server may still have clients still running in a previous release. The default initial JNDI context for EJB clients running on previous versions of WebSphere Application Server is the cell persistent root (legacy root). The home for an enterprise bean deployed in version 5 is bound to its server's server root context. In order for the EJB lookup name for down-level clients to remain unchanged, configure a binding for the EJB home under the cell persistent root.

**Note:** EJB clients running in version 3.5 must be running in version 3.5.5 or above, or in version 3.5.3 or 3.5.4 with e-fix PQ51387 installed.

#### EJB clients running on WebSphere Application Server V5 accessing EJB applications running on WebSphere Application Server V3.5 or V4.0 servers

The default initial context for a WebSphere Application Server v3.5 or v4.0 server is the correct initial context. Simply look up the JNDI name under which the EJB home is bound.

**Note:** To enable WebSphere Application Server V5 clients to access version 3.5.x and 4.0.x servers, the down-level installations must have e-fix PQ60074 installed.

### EJB clients running in an environment other than WebSphere Application Server accessing EJB applications running on WebSphere Application Server V5 servers

When an EJB application running in WebSphere Application Server V5 is accessed by a non-WebSphere Application Server EJB client, the JNDI initial context factory is presumed to be a non-WebSphere Application Server implementation. In this case, the default initial context will be the cell root. If the JNDI service provider being used supports CORBA object URLs, the corbaname format can be used to look up the EJB home. The construction of the stringified name depends on whether the object is installed on a single server or cluster, as shown below.

#### Single server

```
initialContext.lookup("corbaname:iiop:myHost:2809#cell/nodes/node1/servers/server1/myEJB");
```

According to the URL above, the bootstrap host and port are **myHost** and **2809**, and the enterprise bean is installed in a server **server1** in node **node1** and bound in that server under the name **myEJB**.

#### Server cluster

```
initialContext.lookup("corbaname:iiop:myHost:2809#cell/clusters/myCluster/myEJB");
```

According to the URL above, the bootstrap host and port are **myHost** and **2809**, and the enterprise bean is installed in a server cluster named **myCluster** and bound in that cluster under the name **myEJB**.

The above lookup will work with any name server bootstrap host and port configured in the same cell.

The above lookup will also work if the bootstrap host and port belongs to a member of the cluster itself. To avoid a single point of failure, the bootstrap server host and port for each cluster member could be listed in the URL as follows:

```
initialContext.lookup("corbaname:iiop:host1:9810,host2:9810#cell/clusters/myCluster/myEJB");
```

The name prefix **cell/clusters/myCluster/** is not necessary if bootstrapping to the cluster itself, but it will work. The prefix is needed, however, when looking up enterprise beans in other clusters. Name bindings under the **clusters** context are implemented on the name server to resolve to the server root of a running cluster member during a lookup; thus avoiding a single point of failure.

### Without CORBA object URL support

If the JNDI initial context factory being used does not support CORBA object URLs, the initial context can be obtained from the server, and the lookup can be performed on the initial context as follows:

```
Hashtable env = new Hashtable();
env.put(CONTEXT.PROVIDER_URL, "iiop://myHost:2809");
Context ic = new InitialContext(env);
Object o = ic.lookup("cell/clusters/myCluster/myEJB");
```

### Binding resources from MQSeries 5.2

In previous releases of WebSphere Application Server, the MQSeries jmsadmin tool could be used bind resources to the name space. When used with a WebSphere Application Server V5 name space, the resource will be bound within a transient partition in the name space and will not persist past the life of the server process. Instead of binding the MQSeries resources with the jmsadmin tool, bind them from the WebSphere Application Server administrative console, under Resources in the left panel on the console

## JNDI caching

To increase the performance of JNDI operations, the WebSphere Application Server JNDI implementation employs caching to reduce the number of remote calls to the name server for lookup operations. For most cases, use the default cache setting.

When an InitialContext object is instantiated, an association is established between the InitialContext instance and a cache. The initial context and any contexts returned directly or indirectly from a lookup on the initial context are all associated with that same cache instance. By default, the association is based on the provider URL, in particular, the host name and port. The caller can specify the cache name to override this default behavior. A cache instance of a given name is shared by all instances of InitialContext configured to use a cache of that name which were created with the same context class loader in effect. Two EJB applications running in the same server will use their own cache instances, if they are using different context class loaders, even if the cache names are the same.

After an association between an InitialContext instance and cache is established, the association does not change. A `javax.naming.Context` object returned from a lookup operation inherits the cache association of the Context object on which the lookup was performed. Changing cache property values with the `Context.addToEnvironment()` or `Context.removeFromEnvironment()` method does not affect cache behavior. You can change properties affecting a given cache instance with each InitialContext instantiation.

A cache is restricted to a process and does not persist past the life of that process. A cached object is returned from lookup operations until either the max cache life for the cache is reached, or the max entry life for the object's cache entry is reached.

After this time, a lookup on the object causes the cache entry for the object to be refreshed. If a bind or rebind operation is executed on an object, the change is not reflected in any caches other than the one associated with the context from which the bind or rebind was issued. This scenario is most likely to happen when multiple processes are involved, since different processes do not share the same cache, and context objects in all threads in a process typically share the same cache instance for a given name service provider.

Usually, cached objects are relatively static entities, and objects becoming stale are not a problem. However, you can set timeout values on cache entries or on a cache so that cache contents are periodically refreshed.

## JNDI cache settings

Various cache property settings follow. Ensure that all property values are string values.

### **com.ibm.websphere.naming.jndicache.cachename**

The name of the cache to associate with an initial context instance can be specified with this property.

It is possible to create multiple InitialContext instances, each operating on the name space of a different name server. By default, objects from each bootstrap address are cached separately, since they each involve independent name spaces and name collisions could occur if they used the same cache. The provider URL specified when the initial context is created by default serves as the basis for the cache name. With this property, a JNDI client can specify a cache name. Valid options for cache names follow:

| Valid options         | Resulting cache behavior  |
|-----------------------|---|
| providerURL (default) | Use the value for java.naming.provider.url property as the basis for the cache name. Cache names are based on the bootstrap host and port specified in the URL. The bootstrap host is normalized to a fully qualified name, if possible. For example, "corbaname:iiop:server1:2809#some/starting/context" and "corbaloc:iiop://server1" are normalized to the same cache name. If no provider URL is specified, a default cache name is used. |
| Any string            | Use the specified string as the cache name. You can use any arbitrary string with a value other than "providerURL" as a cache name.   |

### **com.ibm.websphere.naming.jndicache.cacheobject**

Turn caching on or off and clear an existing cache with this property.

By default, when an InitialContext is instantiated, it is associated with an existing cache or, if one does not exist, a new one is created. An existing cache is used with its existing contents. In some circumstances, this behavior is not desirable. For example, when objects that are looked up change frequently, they can become stale

in the cache. Other options are available. The following table lists these other options along with the corresponding property value.

| Valid values        | Resulting cache behavior   |
|---------------------|--|
| populated (default) | Use a cache with the specified name. If the cache already exists, leave existing cache entries in the cache; otherwise, create a new cache.  |
| cleared             | Use a cache with the specified name. If the cache already exists, clear all cache entries from the cache; otherwise, create a new cache.   |
| none                | Do not cache. If this option is specified, the cache name is irrelevant. Therefore, this option will not disable a cache that is already associated with other InitialContext instances. The InitialContext that is instantiated is not associated with any cache. |

### **com.ibm.websphere.naming.jndicache.maxcachelife**

Impose a limit to the age of a cache with this property.

By default, cached objects remain in the cache for the life of the process or until cleared with the `com.ibm.websphere.naming.jndicache.cacheobject` property set to "cleared". This property enables a JNDI client to set the maximum life of a cache. This property differs from the `maxentrylife` property (below) in that the entire cache is cleared when the cache lifetime is reached. The table below lists the various `maxcachelife` values and their affect on cache behavior:

| Valid options    | Resulting cache behavior   |
|------------------|--|
| 0 (default)      | Make the cache lifetime unlimited.   |
| Positive integer | Set the maximum lifetime of the entire cache, in minutes, to the specified value. When the maximum lifetime for the cache is reached, the next attempt to read any entry from the cache causes the cache to be cleared |

### **com.ibm.websphere.naming.jndicache.maxentrylife**

Impose a limit to the age of individual cache entries with this property.

By default, cached objects remain in the cache for the life of the process or until cleared with the `com.ibm.websphere.naming.jndicache.cacheobject` property set to cleared. This property enables a JNDI client to set the maximum lifetime of individual cache entries. This property differs from the `maxcachelife` property in that individual entries are refreshed individually as their maximum lifetime reached. This might avoid any noticeable change in performance that might occur if the whole cache is cleared at once. The table below lists the various `maxentrylife` values and their effect on cache behavior:

| Valid options | Resulting cache behavior                |
|---------------|---|
| 0 (default)   | Lifetime of cache entries is unlimited. |

|                  |  |
|------------------|--|
| Positive integer | Set the maximum lifetime of individual cache entries, in minutes, to the specified value. When the maximum lifetime for an entry is reached, the next attempt to read the entry from the cache causes the individual cache entry to refresh. |
|------------------|--|

## Example: Controlling JNDI cache behavior from a program

Following are examples that illustrate how you can use JNDI cache properties to achieve the desired cache behavior. Cache properties take effect when an `InitialContext` object is constructed.

### Usage scenario

```
import java.util.Hashtable;
import javax.naming.InitialContext;
import javax.naming.Context;
import com.ibm.websphere.naming.PROPS;

/*****
Caching discussed in this section pertains to the WebSphere Application Server
initial context factory. Assume the property, java.naming.factory.initial, is set
to "com.ibm.websphere.naming.WsnInitialContextFactory" as a
java.lang.System property.
*****/

Hashtable env;
Context ctx;

// To clear a cache:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_CLEARED);
ctx = new InitialContext(env);

// To set a cache's maximum cache lifetime to 60 minutes:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_MAX_LIFE, "60");
ctx = new InitialContext(env);

// To turn caching off:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_NONE);
ctx = new InitialContext(env);

// To use caching and no caching:

env = new Hashtable();
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_POPULATED);
ctx = new InitialContext(env);
env.put(PROPS.JNDI_CACHE_OBJECT, PROPS.JNDI_CACHE_OBJECT_NONE);
Context noCacheCtx = new InitialContext(env);

Object o;

// Use caching to look up home, since the home should rarely change.
o = ctx.lookup("com/mycom/MyEJBHome");
// Narrow, etc. ...

// Do not use cache if data is volatile.
o = noCacheCtx.lookup("com/mycom/VolatileObject");
// ...
```

## JNDI name syntax

JNDI name syntax is the default syntax and is suitable for typical JNDI clients.

This syntax includes the following special characters: forward slash (/) and backslash (\). Components in a name are delimited by a forward slash. The backslash is used as the escape character. A forward slash is interpreted literally if it is escaped, that is, preceded by a backslash. Similarly, a backslash is interpreted literally if it is escaped.

## INS name syntax

INS syntax is designed for JNDI clients that need to interoperate with CORBA applications.

The INS syntax allows a JNDI client to make the proper mapping to and from a CORBA name. INS syntax is very similar to the JNDI syntax with the additional special character, dot (.). Dots are used to delimit the `id` and `kind` fields in a name component. A dot is interpreted literally when it is escaped. Only one unescaped dot is allowed in a name component. A name component with a non-empty `id` field and empty `kind` field is represented with only the `id` field value and must not end with an unescaped dot. An empty name component (empty `id` and empty `kind` field) is represented with a single unescaped dot. An empty string is not a valid name component representation.

## JNDI to CORBA name mapping considerations

WebSphere Application Server name servers are an implementation of the CORBA CosNaming interface. WebSphere Application Server provides a JNDI implementation which you can use to access CosNaming name servers through the JNDI interface. Issues can exist when mapping JNDI name strings to and from CORBA names.

Each component in a CORBA name consists of an `id` and `kind` field, but a JNDI name component consists of no such fields. Each component in a JNDI name is atomic. Typical JNDI clients do not need to make a distinction between the `id` and `kind` fields of a name component, or know how JNDI name strings map to CORBA names. JNDI clients of this sort can use the JNDI syntax described below. When a name is parsed according to JNDI syntax, each name component is mapped to the `id` field of the corresponding CORBA name component. The `kind` field always has an empty value. This basic syntax is the least obtrusive to the JNDI client in that it has the fewest special characters. However, you cannot represent with this syntax a CORBA name with a non-empty `kind` field. This restriction can prevent EJB applications from interoperating with CORBA applications.

Some clients, however must interoperate with CORBA applications which use CORBA names with non-empty `kind` fields. These JNDI clients must make a distinction between `id` and `kind` so that JNDI names are correctly mapped to CORBA names, particularly when the CORBA names contain components with non-null `kind` fields. Such JNDI clients can use the INS name syntax. With its additional special character, you can use INS to represent any CORBA name. Use of this syntax is not recommended unless it is necessary, because this syntax is more restrictive from the JNDI client's perspective in that the JNDI client must be aware that name components with multiple unescaped dots are syntactically invalid. INS name syntax is part of the OMG CosNaming Interoperable Naming Specification.

## Example: Setting the syntax used to parse name strings

JNDI clients which must interoperate with CORBA applications may need to use INS name syntax to represent names in string format. The name syntax property may be passed to the InitialContext constructor through its parameter, in the System properties, or in a jndi.properties file. The initial context and any contexts looked up from that initial context will parse name strings based on the specified syntax.

The following example shows how to set the name syntax to make the initial context parse name strings according to INS syntax.

### Usage scenario

```
...
import java.util.Hashtable;
import javax.naming.Context;
import javax.naming.InitialContext;
import com.ibm.websphere.naming.PROPS; // WebSphere naming constants
...
Hashtable env = new Hashtable();
env.put(Context.INITIAL_CONTEXT_FACTORY,
        "com.ibm.websphere.naming.WsnInitialContextFactory");
env.put(Context.PROVIDER_URL, ...);
env.put(PROPS.NAME_SYNTAX, PROPS.NAME_SYNTAX_INS);
Context initialContext = new InitialContext(env);
// The following name maps to a CORBA name component as follows:
//   id = "a.name", kind = "in.INS.format"
// The unescaped dot is used as the delimiter.
// Escaped dots are interpreted literally.
java.lang.Object o = initialContext.lookup("a\\.name.in\\.INS\\.format");
...
```

---

## Developing applications that use CosNaming (CORBA Naming interface)

CORBA clients can perform naming operations on WebSphere name servers through the CosNaming interface. The following examples show how to obtain an ORB instance and an initial context as well as how to look up an EJB home.

**Note:** To enable WebSphere Application Server V5 clients to access Versions 3.5.x and 4.0.x servers, the earlier installations must have e-fix PQ60074 installed.

### Steps for this task

1. Get an initial context
2. Perform desired CosNaming operations

## Example: Getting an initial context with CosNaming

In the WebSphere Application Server, an initial context is obtained from a bootstrap server. The address for the bootstrap server consists of a host and port. To get an initial context, you must know the host and port for the server that is used as the bootstrap server.

Obtaining an initial context consists of two basic steps:

1. Obtain an ORB reference
2. Invoke a method on the ORB to obtain the initial reference

These steps are now explained in more detail.



## Obtaining an ORB reference

Pure CosNaming clients, that is clients that are not running in a server process, must create and initialize an ORB instance with which to obtain the initial context. CosNaming clients which run in server processes can obtain a reference to the server ORB with a JNDI lookup. The following examples illustrate how to create and initialize a client ORB and how to obtain a server ORB reference.

### Creating a client ORB instance

To create an ORB instance, invoke the static method, `org.omg.CORBA.ORB.init`. The `init` method requires a property set to the name of the ORB class you want to instantiate. An ORB implementation with the class name `com.ibm.CORBA.iiop.ORB` is included with the WebSphere Application Server. The WebSphere Application Server ORB recognizes additional properties with which you can specify initial references.

The basic steps for creating an ORB are as follows:

1. Create a Properties object.
2. Set the ORB class property to WebSphere Application Server's ORB class.
3. If the bootstrap server is INS-compliant, set the initial reference properties. If the bootstrap server is not INS-compliant (meaning, WebSphere Application Server v4.0.x or earlier), set bootstrap host and port for bootstrap server.
4. Invoke `ORB.init`, passing in the Properties object.

### Usage scenario

```
...
import java.util.Properties;
import org.omg.CORBA.ORB;
...
Properties props = new Properties();
props.put("org.omg.CORBA.ORBClass", "com.ibm.CORBA.iiop.ORB");
props.put("com.ibm.CORBA.ORBInitRef.NameService",
    "corbaloc:iiop:myhost.mycompany.com:2809/NameService");
props.put("com.ibm.CORBA.ORBInitRef.NameServiceServerRoot",
    "corbaloc:iiop:myhost.mycompany.com:2809/NameServiceServerRoot");
// props.put("com.ibm.CORBA.BootstrapHost", "myhost.mycompany.com");
// Use this if bootstrap server is WebSphere 4.0.x or before
// props.put("com.ibm.CORBA.BootstrapPort", "2809");
// Use this if bootstrap server is WebSphere 4.0.x or before
ORB _orb = ORB.init((String[])null, props);
...
```

Notice the initial reference definitions for `NameService` and `NameServiceServerRoot`. The initial context returned for `NameService` depends on the type of bootstrap server. The key `NameServiceServerRoot` is a key introduced in WebSphere Application Server V5. For more information on initial contexts, see the section Initial Contexts.

**Note:** The properties `com.ibm.CORBA.BootstrapHost` and `com.ibm.CORBA.BootstrapPort` are deprecated. They are needed, however, to connect to WebSphere Application Servers of Version 4.0.x or earlier. The default bootstrap host is the local host and the default port is 2809.

### Obtaining a reference to the server ORB

CosNaming clients which run in a server process can obtain a reference to the server ORB with a JNDI lookup on a `java: name`, shown as follows:

### Usage scenario

```
...
import javax.naming.Context;
import javax.naming.InitialContext;
import org.omg.CORBA.ORB;
...
Context initialContext = new InitialContext();
ORB orb = (ORB) initialContext.lookup("java:comp/ORB");
...
```

### Using an ORB reference to get an initial naming reference

There are two basic ways to get an initial CosNaming context. Both ways involve an ORB method invocation. The first way is to invoke the `resolve_initial_references` method on the ORB with an initial reference key. For this call to work, the ORB must be initialized with an initial reference for that key. The other way is to invoke the `string_to_object` method on the ORB, passing in a CORBA object URL with the host and port of the bootstrap server. The following examples illustrate both approaches.

#### Invoking `resolve_initial_references`

Once an ORB reference is obtained, invoke the `resolve_initial_references` method on the ORB to obtain a reference to the initial context. The following code example invokes `resolve_initial_references` on an ORB reference.

### Usage scenario

```
...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Obtain ORB reference as shown in examples earlier in this section
...
org.omg.CORBA.Object obj = _orb.resolve_initial_references("NameService");
NamingContextExt initCtx = NamingContextExtHelper.narrow(obj);
...
```

Note that the key `NameService` is passed to the `resolve_initial_references` method. Other initial context keys are registered in WebSphere Application Servers. For example, `NameServiceServerRoot` can be used to obtain a reference to the server root context in the bootstrap name server. For more information on the initial contexts registered in server ORBs, please see the section `Initial Contexts`.

#### Invoking `string_to_object` with a CORBA object URL

You can use an INS-compliant ORB to obtain an initial context even if the ORB is not initialized with any initial references or bootstrap properties, or if those property settings are for a different server than the name server from which you want to obtain the initial context. To obtain an initial context by explicitly specifying the bootstrap name server, invoke the `string_to_object` method on the ORB, passing in a CORBA object URL which contains the bootstrap server host and port.

The code in the example below invokes the `string_to_object` method on an existing ORB reference, passing in a CORBA object URL which identifies the desired initial context.

### Usage scenario

```

...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Obtain ORB reference as shown in examples earlier in this section
...
org.omg.CORBA.Object obj =
    orb.string_to_object("corbaloc:iiop:myhost.mycompany.com:2809/NameService");
NamingContextExt initCtx = NamingContextExtHelper.narrow(obj);
...

```

Note that the key `NameService` is used in the `corbaloc` URL. Other initial context keys are registered in WebSphere Application Servers. For example, you can use `NameServiceServerRoot` to obtain a reference to the server root context in the bootstrap name server.

### Using an existing ORB and invoking `string_to_object` with a CORBA object URL with multiple name server addresses to get an initial context

CORBA object URLs can contain more than one bootstrap server address. Use this feature when attempting to obtain an initial context from a server cluster. You can specify the bootstrap server addresses for all servers in the cluster in the URL. The operation will succeed if at least one of the servers is running, eliminating a single point of failure. There is no guarantee of any particular order in which the address list will be processed. For example, the second bootstrap server address may be used to obtain the initial context even though the first bootstrap server in the list is available. An example of a `corbaloc` URL with multiple addresses follows.

#### Usage scenario

```

...
import org.omg.CORBA.ORB;
import org.omg.CosNaming.NamingContextExt;
import org.omg.CosNaming.NamingContextExtHelper;
...
// Assume orb is an existing ORB instance org.omg.CORBA.Object obj =
// orb.string_to_object(
"corbaloc::myhost1:9810,:myhost1:9811,:myhost2:9810/NameService");
NamingContextExt initCtx = NamingContextExtHelper.narrow(obj);
...

```

## Example: Looking up an EJB home with `CosNaming`

You can look up an EJB home or other CORBA object from a WebSphere Application Server name server through the CORBA `CosNaming` interface. You can invoke `resolve` or `resolve_str` on the initial context, or you can invoke `string_to_object` on the ORB. You can use a qualified name so that the name resolves regardless of which name server the lookup is executed on, or use an unqualified name that only resolves from the server root context on the name server that actually contains the object binding. (The qualified name traverses the federated system name space to the specified server root context.)

### Qualified and unqualified names

Each application server contains a name server. System artifacts such as EJB homes are bound in that name server. The various name servers are federated by means of a system name space structure. The recommended way to look up objects on different servers is to use a qualified name. A qualified name can be a topology-based name, based on the name of the cluster or single server and node that contains the object. You can define fixed qualified names for objects. With

qualified names, you can look up objects residing on different servers from the same initial context by traversing the system name space structure. Alternatively, you can use an unqualified name, but an unqualified name will only resolve using the name server associated with the object's application server.

### **CosNaming.resolve (and resolve\_str) vs. ORB.string\_to\_object**

If you have an initial context from any name server in a WebSphere Application Server cell, you can look up any CORBA object with a qualified name. You do not need additional host and port information for the target object's name server.

Alternatively, you can look up an object by invoking `string_to_object` on the ORB, passing in a corbaname URL. Typically, an IIOP type URL is specified, so the bootstrap address information required for an initial context must be contained in the URL. You can use a qualified or unqualified stringified name, but an unqualified name resolves only if the initial context is from the name server in which the object is bound.

The following examples show CosNaming resolve operations using qualified topology-based lookup names and an unqualified lookup name.

### **CosNaming resolve operation using a qualified name**

The topology-based qualified name for an object depends on whether the object is bound in a single server or a server cluster. Examples of each follow.

#### **Single Server**

The following example shows the lookup of an EJB home that is running in a single server. The enterprise bean that is being looked up is running in the server, `MyServer`, on the node, `Node1`.

```
// Get the initial context as shown in the previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the name under which the EJB home is bound
org.omg.CORBA.Object ejbHome = initialContext.resolve_str(
    "cell/nodes/Node1/servers/MyServer/mycompany/accounting/AccountEJB");
accountHome =
    (AccountHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, AccountHome.class);
```

#### **Server Cluster**

The following example shows a lookup of an EJB home that is running in a cluster. The enterprise bean being that is looked up is running in the cluster, `Cluster1`. The name can be resolved if any of the cluster members is running.

#### **Usage scenario**

```
// Get the initial context as shown in the previous example
// Using the form of lookup name below, it doesn't matter which
// server in the cell is used to obtain the initial context.
...
// Look up the home interface using the name under which the EJB home is bound
org.omg.CORBA.Object ejbHome = initialContext.resolve_str(
    "cell/clusters/Cluster1/mycompany/accounting/AccountEJB");
accountHome =
    (AccountHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, AccountHome.class);
```

## ORB string\_to\_object operation using an unqualified stringified name

If the resolve operation is being performed on the name server that contains the object, the system name space does not need to be traversed, and you can use an unqualified lookup name. Note that this name does not resolve on other name servers. If an unqualified name is provided, the object key must be `NameServiceServerRoot` so that the correct initial context is selected. If a qualified name is provided, you can use the default key of `NameService`.

The following example shows a lookup of an EJB home. The enterprise bean that is being looked up is bound on the name server running on the host `myHost` on port 2809. Note the object key of `NameServiceServerRoot`.

### Usage scenario

```
// Assume orb is an existing ORB instance
...
// Look up the home interface using the name under which the EJB home is bound
org.omg.CORBA.Object ejbHome = orb.string_to_object(
    "corbaname:iiop:myHost:2809/NameServiceServerRoot#mycompany/accounting");
AccountHome =
    (AccountHome)javax.rmi.PortableRemoteObject.narrow(ejbHome, AccountHome.class);
```

---

## Configured name bindings

Administrators can configure bindings into the name space. A configured binding is different from a programmatic binding in that the system creates the binding every time a server is started, even if the target context is in a transient partition.

Administrators can add name bindings to the name space through the configuration. Name servers add these configured bindings to the name space view, by reading the configuration data for the bindings. Configuring bindings is an alternative to creating the bindings from a program. Configured bindings have the advantage of being created each time a server starts, even when the binding is created in a transient partition of the name space. Cell-scoped configured bindings provide interoperability with JNDI clients running on previous versions of WebSphere Application Server. Additionally, you can configure cell-scoped bindings to create a fixed qualified name for server objects.

### Scope

You can configure a binding at one of the following three scopes: cell, node, or server. Cell-scoped bindings are created under the cell persistent root context. Node-scoped bindings are created under the node persistent root context for the specified node. Server-scoped bindings are created under the server root context for the selected server. If the target server of a server-scoped binding is a cluster, the binding is created under the server root context of each cluster member.

**Note:** The term *server* includes clusters and can be used interchangeably with the term *cluster* with respect to configured bindings. When applied to a cluster, a server-scoped binding is created in the server root for all member servers.

The scope you select for new bindings depends on how the binding is to be used. For example, if the binding is not specific to any particular node or server, or if you do not want the binding to be associated with any specific node or server, a cell-scoped binding is a suitable scope. Defining fixed names for enterprise beans to create fixed qualified names is just such an application. If a binding is to be used only by clients of an application running on a particular server, or if you

want to configure a binding with the same name on different servers which resolve to different objects, a server-scoped binding would be appropriate. Note that two servers can have configured bindings with the same name but resolve to different objects. At the cell scope, only one binding with a given name can exist.

### **Intermediate Contexts**

Intermediate contexts created with configured bindings are read-only. For example, if an EJB home binding is configured with the name `some/compound/name/ejbHome`, the intermediate contexts `some`, `some/compound`, and `some/compound/name` will be created as read-only contexts. You cannot add, update, or remove any read-only bindings.

The configured binding name cannot conflict with existing bindings. However, configured bindings can use the same intermediate context names. Therefore, a configured binding with the name `some/compound/name2/ejbHome2` does not conflict with the previous example name.

## **Configured binding types**

Types of objects that you can bind follow:

### **EJB: EJB home installed in some server in the cell**

The following data is required to configure an EJB home binding:

- JNDI name of the EJB server or server cluster where the enterprise bean is deployed
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.

This type of binding is of special significance because you can use it to provide interoperability with WebSphere Application Server v3.5.x and v4.0.x JNDI clients. The default initial context for these earlier clients is the cell persistent root, which is different from the initial context of the server root for WebSphere Application Server V5 JNDI clients. If you migrate an application to the current release, you can configure an EJB binding at the cell scope so that the lookup names for the enterprise bean do not change for clients still running in an earlier WebSphere Application Server version.

A cell-scoped EJB binding is also useful for creating a fixed lookup name for an enterprise bean so that the qualified name is not dependent on the topology.

### **CORBA: CORBA object available from some CosNaming name server**

You can identify any CORBA object bound into some INS compliant CosNaming server with a `corbaname` URL. The referenced object does not have to be available until the binding is actually referenced by some application.

The following data is required in order to configure a CORBA object binding:

- The `corbaname` URL of the CORBA object
- An indicator if the bound object is a context or leaf node object (to set the correct CORBA binding type of context or object).
- Target root for the configured binding
- The name of the configured binding, relative to the target root.

**Indirect: Any object bound in WebSphere Application Server name space accessible with JNDI**

Besides CORBA objects, this includes `javax.naming.Referenceable`, `javax.naming.Reference`, and `java.io.Serializable` objects. The target object itself is not bound to the name space. Only the information required to look up the object is bound. Therefore, the referenced name server does not have to be running until the binding is actually referenced by some application. The following data is required in order to configure an indirect JNDI lookup binding:

- JNDI provider URL of name server where object resides
- JNDI lookup name of object
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.

A cell-scoped indirect binding is useful when creating a fixed lookup name for a resource so that the qualified name is not dependent on the topology. You can also achieve this topology by widening the scope of the resource definition.

**Note:** WebSphere Application Server v3.5.x clients cannot access this type of binding .

**String: String constant**

You can configure a binding of a string constant. The following data is required to configure a string constant binding:

- String constant value
- Target root for the configured binding (scope)
- The name of the configured binding, relative to the target root.

---

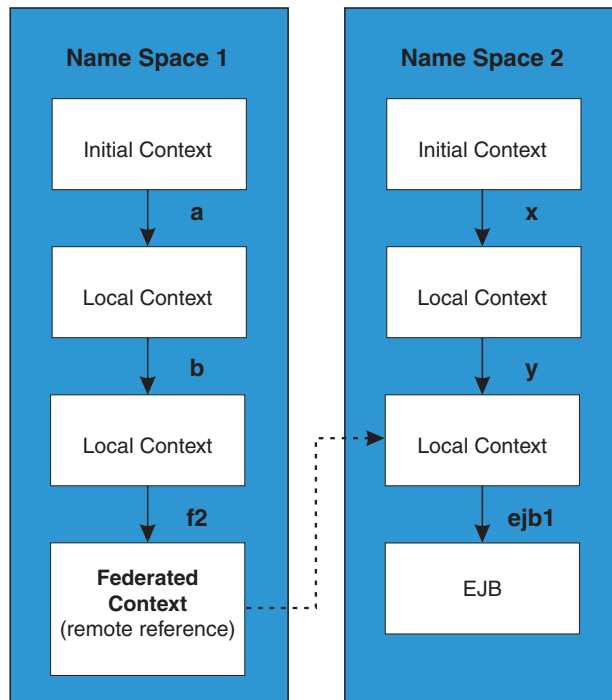
## Name space federation

Federating name spaces involves binding contexts from one name space into another name space.

For example, assume that a name space, Name Space 1, contains a context under the name `a/b`. Also assume that a second name space, Name Space 2, contains a context under the name `x/y`. (See the following illustration.) If context `x/y` in Name Space 2 is bound into context `a/b` in Name Space 1 under the name `f2`, the two name spaces are federated. Binding `f2` is a federated binding because the context associated with that binding comes from another name space. From Name Space 1, a lookup of the name `a/b/f2` returns the context bound under the name `x/y` in Name Space 2. Furthermore, if context `x/y` contains an Enterprise JavaBeans (EJB) home bound under the name `ejb1`, the EJB home could be looked up from Name Space 1 with the lookup name `a/b/f2/ejb1`. Notice that the name crosses name

spaces. This fact is transparent to the naming client.

### Federated Name Spaces



In a WebSphere Application Server name space, you can create federated bindings with the following restrictions:

- Federation is limited to CosNaming name servers. A WebSphere Application Server name server is a Common Object Request Broker Architecture (CORBA) CosNaming implementation. You can create federated bindings to other CosNaming contexts. You cannot, for example, bind contexts from an LDAP name server implementation.
- If you use JNDI to federate the name space, you must use WebSphere Application Server's initial context factory to obtain the reference to the federated context. If you use some other initial context factory implementation, you either may not be able to create the binding, or the level of transparency may be reduced.
- A federated binding to a non-WebSphere Application Server naming context has the following functional limitations:
  - JNDI operations are restricted to the use of CORBA objects. For example, you can look up EJB homes, but you cannot look up non-CORBA objects such as data sources.
  - JNDI caching is not supported for non-WebSphere Application Server name spaces. This restriction affects the performance of lookup operations only.
- Do not federate two WebSphere Application Server standalone server name spaces. Incorrect behavior may result. If you want to federate WebSphere Application Server name spaces, you should use servers running under the Network Deployment or Enterprise packages of WebSphere Application Server.



---

## Name space bindings

Administrators can add name bindings to the name space through the configuration. Name servers add these configured bindings to the name space view by reading the configuration data for the bindings. Configuring bindings is an alternative to creating the bindings from a program.

Configured bindings are created each time a server starts, even when the binding is created in a transient partition of the name space. One major use of configured bindings is to provide interoperability with JNDI clients running on previous versions of the WebSphere Application Server.

There are four different kinds of bindings that you can configure:

- Enterprise JavaBeans (EJB)
- CORBA object
- Indirect Lookup
- String

---

## Configuring and viewing name space bindings

To view or configure an EJB, CORBA, Indirect lookup or string name space binding, complete the following:

### Steps for this task

1. Open the administrative console.
2. Click **Environment**.
3. Click **Manage Name Space Bindings**.
4. Select the desired scope by entering in a node name for node-scoped bindings, or a node name and server name for server-scoped bindings, and click **Apply**.
5. To create a new binding, click **New** and follow the instructions. To edit a previously created binding, click the binding you want to edit and proceed to the next step.
6. Edit the **Binding identifier**, the **Name in name space**, and the **String value** fields as desired.  
**Note:** All of these fields are required.
7. Click **Finish** to register the changes.

## String binding settings

Use this page to configure a new string binding or to view or edit an existing string binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > *string\_namespace\_binding***.

### **Scope**

Shows the scope of the configured binding. This value indicates the configuration location for the `namebindings.xml` file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

### **Binding Type**

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

### **Binding Identifier**

Specifies the name that uniquely identifies this configured binding.

### **Name in Name Space**

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

### **String Value**

Specifies the string to be bound into the name space.

## **CORBA object binding settings**

Use this page to configure a new name binding of a CORBA object binding, or to view or edit an existing CORBA object binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > CORBA\_namespace\_binding**.

### **Scope**

Shows the scope of the configured binding. This value indicates the configuration location for the namebindings.xml file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

### **Binding Type**

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

### **Binding Identifier**

Specifies the name that uniquely identifies this configured binding.

### **Name in Name Space**

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

### **Corbaname URL**

Specifies the CORBA name URL string identifying where the object is bound in a CosNaming server.

### **Federated Context**

Specifies whether the target is a CosNaming context (true) or a leaf node object (false).

| Value | Result  |
|-------|---|
| true  | The target object is bound with a context CORBA binding type. If the corbaname URL does not resolve to a NamingContext, an error occurs when the binding is first used (which is when the URL is first resolved). |
| false | The target object is bound with an object CORBA binding type.   |

## Indirect lookup binding settings

Use this page to configure a new indirect lookup name binding, or to view or edit an existing indirect lookup binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > *indirect\_lookup\_namespace\_binding***.

### Scope

Shows the scope of the configured binding. This value indicates the configuration location for the `namebindings.xml` file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

### Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

### Binding Identifier

Specifies the name that uniquely identifies this configured binding.

### Name in Name Space

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

### Provider URL

Specifies the provider URL string needed to obtain a JNDI initial context.

### JNDI Name

Specifies the name used to look up the target object from the initial context.

## EJB binding settings

Use this page to configure a new EJB binding, or to view or edit an existing EJB binding.

To view this administrative console page, click **Environment > Naming > Name Space Bindings > *EJB\_namespace\_binding***.

### Scope

Shows the scope of the configured binding. This value indicates the configuration location for the `namebindings.xml` file. This field is for information purposes only and cannot be updated.

If the configured binding is cell-scoped, the starting context is the cell persistent root context. If the configured binding is node-scoped, the starting context is the node persistent root context. If the configured binding is server-scoped, the starting context is the server's server root context.

### Binding Type

Shows the type of binding configured. Possible choices are String, EJB, CORBA, and Indirect. This field is for information purposes only and cannot be updated.

### **Binding Identifier**

Specifies the name that uniquely identifies this configured binding.

### **Name in Name Space**

Specifies the name used for this binding in the name space. This name can be a simple or compound name depending on the portion of the name space where this binding is configured.

### **Enterprise Bean Location**

Specifies whether the enterprise bean is running in a server cluster or a single server. If Single Server is specified, type the node name.

### **Server**

Specifies the name of the cluster or non-clustered server in which the enterprise bean is configured.

### **JNDI Name**

Specifies the JNDI name of the deployed enterprise bean (the bean's JNDI name that is in the enterprise bean bindings—not the java:comp name)

## **Name space binding collection**

Use this page to configure a name binding of an EJB, a CORBA CosNaming NamingContext, a CORBA leaf node object, an object that you can look up using JNDI, or a constant string value.

Binding information for configured bindings is stored in the configuration and applied upon startup of the name server for each server within the scope of the binding.

To view the Manage Name Space Bindings Settings page, Click **Environment > Naming > Name Space Bindings**.

Click the check boxes to select one or more of the users in your collection. Use the buttons to control the selected users.

### **Name space bindings**

Shows the names given to uniquely identify these configured bindings.

---

## **Configuring name servers**

To configure a name server, complete the following:

### Steps for this task

1. Open the administrative console.
2. Click **Servers**.
3. Click **Application Servers**.
4. Click the application server you want to configure.
5. Click **Server Components**.
6. Click **Name Server**.
7. Edit the fields as desired.  
**Note:** All of these fields are mandatory.
8. **(Optional)** To make other changes, click **Custom Properties**.
9. Click **OK** to register your changes.

## Name server settings

Use this page to configure Naming Service Provider settings for the application server.

To view this administrative console page, click one of the following paths:

- **Servers > Application Servers > *server\_name* > Server Components > Name Server**
- **Servers > JMS Servers > *server\_name* > Server Components > Name Server**

The **Configuration** tab on the page provides editable fields and the **Runtime** tab provides read-only information.

### Name

Specifies the display name for the server.

### Initial State

Specifies the execution state.

---

## Troubleshooting name space problems

Many naming problems can be avoided by fully understanding the key underlying concepts of WebSphere Application Server naming.

### Steps for this task

1. Review the key concepts of WebSphere Application Server naming, especially Name space logical view and Lookup names support in deployment descriptors and thin clients.
2. Review the programming examples that are included in the sections explaining the JNDI and CosNaming interfaces.
3. Read "Naming services component troubleshooting tips" (not in this document) for additional general information.
4. If you "Cannot look up an object hosted by WebSphere Application Server from a servlet, JSP file, or other client" (not in this document), read this article.

## dumpNameSpace tool

You can use the name space dump utility to dump the contents of a name space accessed through a name server. When you invoke the dump utility, the naming service must be active. The dump utility cannot dump name spaces local to the server process, such as those with java: and local: URL schemes. The local: name space contains references to enterprise beans with local interfaces. Use the name space dump utility for java:, local: and server name spaces to dump java: and local: name spaces.

Note that the server root context for the server at the specified host and port is dumped (unless a non-default starting context which precludes it is specified). The server root contexts for other servers are not dumped.

Command line invocation descriptions of the name space dump utility follow. This section includes sample output.

You can also access this utility through its program interface. Refer to the class `com.ibm.websphere.naming.DumpNameSpace` in the WebSphere Application Server API documentation.

## Syntax

To invoke the tool through the command line, enter the following command from the *WebSphere/AppServer/bin* directory:

| Platform   | Command                                |
|------------|--|
| UNIX       | dumpNameSpace.sh [[-keyword value]...] |
| Windows NT | dumpNameSpace [[-keyword value]...]    |

## Parameters

The keywords and associated values for the dumpNameSpace utility follow:

**-host *myhost.austin.ibm.com***

Indicates the bootstrap host or the WebSphere Application Server host whose name space you want to dump. The value defaults to *localhost*.

**-port *nnn***

Indicates the bootstrap port which, if not specified, defaults to **2809**.

**-root {cell | server | node | host | legacy | tree | default}**

Indicates the root context to use as the initial context for the dump. The applicable root options and default root context depend on the type of name server from which the dump is being obtained. This information is provided in the following tables.

For WebSphere Application Servers V5 or later:

|        |  |
|--------|--|
| cell   | DumpNameSpace default. Dump the tree starting at the cell root context.  |
| server | Dump the tree starting at the server root context.                       |
| node   | Dump the tree starting at the node root context. (Synonymous with host.) |

For WebSphere Application Servers v4.0 or later:

|        |  |
|--------|--|
| legacy | DumpNameSpace default. Dump the tree starting at the legacy root context.          |
| host   | Dump the tree starting at the bootstrap host root context. (Synonymous with node.) |
| tree   | Dump the tree starting at the tree root context.                                   |

For all WebSphere Application Servers and other name servers:

|         |   |
|---------|---|
| default | Dump the tree starting at the initial context which JNDI returns by default for that server type. This is the only -root choice that is compatible with WebSphere Application Servers prior to v4.0 and with non-WebSphere Application Server name servers. |
|---------|---|

**-url *some provider URL***

Indicates the value for the java.naming.provider.url property used to get

the initial JNDI context. This option can be used in place of the -host, -port, and -root options. If the -url option is specified, the -host, -port, and -root options are ignored.

**-factory** *com.ibm.websphere.naming.WsnInitialContextFactory*

Indicates the initial context factory to be used to get the JNDI initial context. The value defaults to:

*com.ibm.websphere.naming.WsnInitialContextFactory* The default value generally does not need to be changed.

**-startAt** *some/subcontext/in/the/tree*

Indicates the path from the bootstrap host's root context to the top level context where the dump should begin. The utility recursively dumps subcontexts below this point. It defaults to an empty string, that is, the bootstrap host root context.

**-format** {jndi | ins}

| Option | Description  |
|--------|--|
| jndi   | The default. Displays name components as atomic strings. |
| ins    | Shows name components parsed per INS rules (id.kind).    |

**-report** {short | long}

| Option | Description  |
|--------|--|
| short  | The default. Dumps the binding name and bound object type. This output is also provided by JNDI Context.list().  |
| long   | <p>Dumps the binding name, bound object type, local object type, and string representation of the local object (that is, the IORs, string values, and other values that are printed).</p> <p>For objects of user-defined classes to display correctly with the long report option, it may be necessary to add their containing directories to the list of directories searched. Set the environment variable <b>WAS_USER_DIRS</b>. The value can include one or more directories, as for example:</p> <p><b>Platform</b><br/>           Command</p> <p><b>UNIX</b>   WAS_USER_DIRS=/usr/classdir1:/usr/classdir2<br/>           export WAS_USER_DIRS</p> <p><b>Windows NT</b><br/>           set WAS_USER_DIRS=c:\classdir1;d:\classdir2</p> <p>All zip, jar, and class files in the specified directories can then be resolved by the class loader when running <b>dumpNameSpace</b>.</p> |

**-traceString** *"some.package.name.to.trace.\*=all=enabled"*

Represents the trace string with the same format as that generated by the servers. The output is sent to the file, DumpNameSpaceTrace.out.

## Example: Invoking the name space dump utility

It is often helpful to view a dump of the name space to understand why a naming operation is failing. You can invoke the name space dump utility from the command line or from a program. Examples of each option follow.

### Invoking name space dump utility from a command line

Invoke the name space dump utility from the command line by entering the following command:

```
dumpNameSpace -host myhost.mycompany.com -port 901
```

OR

```
dumpNameSpace -url corbaloc:iiop:myhost.mycompany.com:901
```

There are several command line options to choose from. For detailed help, enter the following command:

```
dumpNameSpace -help
```

### Invoking name space dump utility from a Java program

You can dump name spaces from a program with the `com.ibm.websphere.naming.DumpNameSpace` API. Refer to the WebSphere Application Server API documentation for details on the `DumpNameSpace` program interface.

The following example illustrates how to invoke the name space dump utility from a Java program:

```
{
...
import javax.naming.Context;
import javax.naming.InitialContext;
import com.ibm.websphere.naming.DumpNameSpace;
...
java.io.PrintStream filePrintStream = ...
Context ctx = new InitialContext();
ctx = (Context) ctx.lookup("cell/nodes/node1/servers/server1"); // Starting context for dump
DumpNameSpace dumpUtil = new DumpNameSpace(filePrintStream, DumpNameSpace.SHORT);
dumpUtil.generateDump(ctx);
...
}
```

## Name space dump utility for java:, local: and server name spaces

Sometimes it is helpful to dump the java: name space for a J2EE application. You cannot use the `dumpNameSpace` command line utility for this purpose because the application's java: name space is accessible only by that J2EE application. From the WebSphere Application Server scripting tool, you can invoke a `NameServer` MBean to dump the java: name space for any J2EE application running in that same server process.

There is another name space local to server process which you cannot dump with the `dumpNameSpace` command line utility. This name space has the URL scheme of `local:` and is used by the container to bind objects locally instead of through the name server. The `local:` name space contains references to enterprise beans with local interfaces. There is only one `local:` name space in a server process. You can dump the `local:` name space by invoking the `NameServer` MBean associated with that server process.



## Name space dump options

Name space dump options are specified in the MBean invocation as a parameter in character string format. The option descriptions follow.

### **-startAt** *some/subcontext/in/the/tree*

Indicates the path from the name space root context to the top level context where the dump should begin. The utility recursively dumps subcontexts below this point. It defaults to an empty string, that is, the root context.

### **-report** {short | long}

| Option | Description  |
|--------|--|
| short  | The default. Dumps the binding name and bound object type. This output is also provided by JNDI Context.list().  |
| long   | Dumps the binding name, bound object type, local object type, and string representation of the local object (that is, the IORs, string values, and other values that are printed). |

### **-root** {tree | host | legacy | cell | node | server | | default}

Specify the root context of where the dump should start. The default value for -root is *cell*. This option is only valid for server name space dumps.

| Option  | Description  |
|---------|--|
| tree    | Dumps the tree starting at the tree root context.  |
| host    | Dumps the tree starting at the server host root context (synonymous with "node").                  |
| legacy  | Dumps the tree starting at the legacy root context.  |
| cell    | The default. Dumps the tree starting at the cell root context.                                     |
| node    | Dumps the tree starting at the node root context (synonymous with "host").                         |
| server  | Dumps the tree starting at the server root context. This is -root default.                         |
| default | Dumps the tree starting at the initial context which JNDI returns by default for that server type. |

### **-format** {jndi | ins}

Specify the format to display name component as atomic strings or parsed according to INS rules (id.kind). This option is only valid for server name space dumps.

| Option | Description  |
|--------|--|
| jndi   | Displays name components as atomic strings. This is -format default. |
| ins    | Displays name components parsed according to INS rules (id.kind).    |

## NameServer MBean invocation

**Enter the WebSphere Application Server scripting command prompt:** Invoke a method on a NameServer MBean by using the WebSphere Application Server scripting tool. Enter the scripting command prompt by typing the following command:

| Platform | Command    |
|----------|------------|
| UNIX     | wsadmin.sh |

| Platform   | Command |
|------------|---------|
| Windows NT | wsadmin |

Use the `-help` option for help on using the `wsadmin` command.

**Select the NameServer MBean instance to invoke:** Execute the following script commands to select the NameServer instance you want to invoke. For example (shown here on multiple lines for publication),

```
set mbean [$AdminControl completeObjectName
  WebSphere:*,type=NameServer,cell=cellName,
  node=nodeName,process=serverName]
```

where *cellName*, *nodeName*, and *serverName* are the names of the cell, node, and server for the MBean you want to invoke. The specified server must be running before you can invoke a method on the MBean.

You can see a list of all NameServer MBeans current running by issuing the following query:

```
$AdminControl queryNames {*:*,type=NameServer}
```

#### Invoke the NameServer MBean: `java: name space`

Dump a `java: name space` by invoking the `dumpJavaNameSpace` method on the NameServer MBean. Since each server application has its own `java: name space`, the application must be specified on the method invocation. An application is identified by the application name, module name, and component name. The method syntax follows:

```
$AdminControl invoke $mbean dumpJavaNameSpace {{appName}{modName}{compName}{opts}}
```

where *appName* is the application name, *modName* is the module name, and *compName* is the component name of the `java: name space` you want to dump. The value for *opts* is the list of name space dump options described earlier in this section. The list can be empty.

#### local: name space

Dump a `java: name space` by invoking the `dumpLocalNameSpace` method on the NameServer MBean. Since there is only one local: name space in a server process, you have to specify the name space dump options only.

```
$AdminControl invoke $mbean dumpLocalNameSpace {{opts}}
```

where *opts* is the list of name space dump options described earlier in this section. The list can be empty.

#### Name space dump output

Name space dump output is sent to the console. It is also written to the file `DumpNameSpace.log`, in the server's log directory.

### Name space dump sample output

Name space dump output looks like the following example, which is the **SHORT** dump format:

```
Getting the initial context
Getting the starting context
```

```

=====
Name Space Dump
  Provider URL: corbaloc:iiop:localhost:9810
  Context factory: com.ibm.websphere.naming.WsnInitialContextFactory
  Requested root context: cell
  Starting context: (top)=outpostNetwork
  Formatting rules: jndi
  Time of dump: Mon Sep 16 18:35:03 CDT 2002
=====

```

```

=====
Beginning of Name Space Dump
=====

```

```

  1 (top)
  2 (top)/domain                javax.naming.Context
  2   Linked to context: outpostNetwork
  3 (top)/cells                 javax.naming.Context
  4 (top)/clusters              javax.naming.Context
  5 (top)/clusters/Cluster1     javax.naming.Context
  6 (top)/cellname              java.lang.String
  7 (top)/cell                  javax.naming.Context
  7   Linked to context: outpostNetwork
  8 (top)/deploymentManager      javax.naming.Context
  8   Linked to URL: corbaloc::outpost:9809/NameServiceServerRoot
  9 (top)/nodes                 javax.naming.Context
 10 (top)/nodes/will2           javax.naming.Context
 11 (top)/nodes/will2/persistent javax.naming.Context
 12 (top)/nodes/will2/persistent/SomeObject SomeClass
 13 (top)/nodes/will2/nodename  java.lang.String
 14 (top)/nodes/will2/domain    javax.naming.Context
 14   Linked to context: outpostNetwork
 15 (top)/nodes/will2/cell      javax.naming.Context
 15   Linked to context: outpostNetwork
 16 (top)/nodes/will2/servers   javax.naming.Context
 17 (top)/nodes/will2/servers/server1 javax.naming.Context
 18 (top)/nodes/will2/servers/will2 javax.naming.Context
 19 (top)/nodes/will2/servers/member2 javax.naming.Context
 20 (top)/nodes/will2/node      javax.naming.Context
 20   Linked to context: outpostNetwork/nodes/will2
 21 (top)/nodes/will2/nodeAgent javax.naming.Context
 22 (top)/nodes/outpost         javax.naming.Context
 23 (top)/nodes/outpost/node    javax.naming.Context
 23   Linked to context: outpostNetwork/nodes/outpost
 24 (top)/nodes/outpost/nodeAgent javax.naming.Context
 24   Linked to URL: corbaloc::outpost:2809/NameServiceServerRoot
 25 (top)/nodes/outpost/persistent javax.naming.Context
 26 (top)/nodes/outpost/nodename java.lang.String
 27 (top)/nodes/outpost/domain  javax.naming.Context
 27   Linked to context: outpostNetwork
 28 (top)/nodes/outpost/servers javax.naming.Context
 29 (top)/nodes/outpost/servers/server1 javax.naming.Context
 30 (top)/nodes/outpost/servers/server1/url javax.naming.Context
 31 (top)/nodes/outpost/servers/server1/url/CatalogDAOURL
 31   java.net.URL
 32 (top)/nodes/outpost/servers/server1/mail javax.naming.Context
 33 (top)/nodes/outpost/servers/server1/mail/PlantsByWebSphere
 33   javax.mail.Session
 34 (top)/nodes/outpost/servers/server1/TransactionFactory
 34   com.ibm.ejs.jts.jts.ControlSet$LocalFactory
 35 (top)/nodes/outpost/servers/server1/servername java.lang.String
 36 (top)/nodes/outpost/servers/server1/WSsamples javax.naming.Context
 37 (top)/nodes/outpost/servers/server1/WSsamples/TechSampDataSource
 37   TechSamp
 38 (top)/nodes/outpost/servers/server1/thisNode javax.naming.Context
 38   Linked to context: outpostNetwork/nodes/outpost
 39 (top)/nodes/outpost/servers/server1/cell javax.naming.Context
 39   Linked to context: outpostNetwork

```

```

40 (top)/nodes/outpost/servers/server1/eis          javax.naming.Context
41 (top)/nodes/outpost/servers/server1/eis/DefaultDataSource_CMP
41                                     Default_CF
42 (top)/nodes/outpost/servers/server1/eis/WSSamples javax.naming.Context
43 (top)/nodes/outpost/servers/server1/eis/WSSamples/TechSampDataSource_CMP
43                                     TechSamp_CF
44 (top)/nodes/outpost/servers/server1/eis/jdbc      javax.naming.Context
45 (top)/nodes/outpost/servers/server1/eis/jdbc/PlantsByWebSphereDataSource_CMP
45                                     PLANTSDB_CF
46 (top)/nodes/outpost/servers/server1/eis/jdbc/petstore
46                                     javax.naming.Context
47 (top)/nodes/outpost/servers/server1/eis/jdbc/petstore/PetStoreDB_CMP
47                                     PetStore_CF
48 (top)/nodes/outpost/servers/server1/eis/jdbc/CatalogDB_CMP
48                                     Catalog_CF
49 (top)/nodes/outpost/servers/server1/jta          javax.naming.Context
50 (top)/nodes/outpost/servers/server1/jta/usertransaction
50                                     java.lang.Object
51 (top)/nodes/outpost/servers/server1/DefaultDataSource
51                                     Default DataSource
52 (top)/nodes/outpost/servers/server1/jdbc          javax.naming.Context
53 (top)/nodes/outpost/servers/server1/jdbc/CatalogDB CatalogDB
54 (top)/nodes/outpost/servers/server1/jdbc/petstore javax.naming.Context
55 (top)/nodes/outpost/servers/server1/jdbc/petstore/PetStoreDB
55                                     PetStoreDB
56 (top)/nodes/outpost/servers/server1/jdbc/PlantsByWebSphereDataSource
56                                     PLANTSDB
57 (top)/nodes/outpost/servers/outpost              javax.naming.Context
57   Linked to URL: corbaloc::outpost:2809/NameServiceServerRoot
58 (top)/nodes/outpost/servers/member1             javax.naming.Context
59 (top)/nodes/outpost/cell                        javax.naming.Context
59   Linked to context: outpostNetwork
60 (top)/nodes/outpostManager                     javax.naming.Context
61 (top)/nodes/outpostManager/domain              javax.naming.Context
61   Linked to context: outpostNetwork
62 (top)/nodes/outpostManager/cell                javax.naming.Context
62   Linked to context: outpostNetwork
63 (top)/nodes/outpostManager/servers             javax.naming.Context
64 (top)/nodes/outpostManager/servers/dmgr        javax.naming.Context
64   Linked to URL: corbaloc::outpost:9809/NameServiceServerRoot
65 (top)/nodes/outpostManager/node                javax.naming.Context
65   Linked to context: outpostNetwork/nodes/outpostManager
66 (top)/nodes/outpostManager/nodename            java.lang.String
67 (top)/persistent                              javax.naming.Context
68 (top)/persistent/cell                         javax.naming.Context
68   Linked to context: outpostNetwork
69 (top)/legacyRoot                              javax.naming.Context
69   Linked to context: outpostNetwork/persistent
70 (top)/persistent/AnotherObject                 AnotherClass

```

```

=====
End of Name Space Dump
=====

```


---

## Naming and directories: Resources for learning

Use the following links to find relevant supplemental information about naming and directories. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

### Programming specifications

-  **OMG CosNaming Interoperable Naming Specification**  
(<http://java.sun.com/j2se/1.4.1/docs/api/org/omg/CosNaming/package-summary.html>)



---

## Chapter 25. Using the dynamic cache service to improve performance

The dynamic cache service works within an application server Java virtual machine (JVM), intercepting calls to cacheable objects. For example, it intercepts calls through a servlet `service()` method or a command `execute()` method, and either stores the output of the object to, or serves the content of the object from the dynamic cache.

WebSphere Application Server, Version 4.0, supported the configuration of dynamic servlet caching through the `servletcache.xml` file. To utilize the new and improved functionality of the dynamic cache service, configure your cache policy using the `cachespec.xml` format.

The dynamic caching documentation provides you with the following tasks to enable and configure the dynamic cache service, as well as advanced features, such as controlling external caches and building user-defined drop-in components to customize the cache operation.

### Steps for this task

1. Enable the dynamic cache service globally.
2. Configure servlet caching.
3. Configure Edge Side Include (ESI) caching.
4. Configure command caching.
5. Configure Web services caching.

### What to do next

To use the `DistributedMap` interface for the dynamic cache, see [Using the DistributedMap interface for the dynamic cache](#).

---

## Dynamic cache

Caching the output of servlets, commands and JavaServer Pages (JSP) files, improves application performance. WebSphere Application Server consolidates several caching activities, including servlets, Web services, and WebSphere commands into one service called the *dynamic cache*. These caching activities work together to improve application performance, and share many configuration parameters, which are set in the dynamic cache service of an application server.

You can use the dynamic cache to improve the performance of servlet and JSP files by serving requests from an in-memory cache. Cache entries contain servlet output, results of servlet execution, and metadata.

---

## Configuring cache replication

Cache replication leverages the WebSphere internal replication service that is also leveraged for HttpSession memory-to-memory replication for failover purposes. Hence, a replication domain with at least one replicator entry needs to exist in order to replicate the data. The dynamic cache service, in essence, connects to the replicator. See more information in the topic referring to managing internal replication.

To configure cache replication and its features:

### Steps for this task

1. Click **Servers>Manage Application Servers** in the administrative console navigation tree.
2. Click *server*.
3. Click **Dynamic Cache Service**.
4. Click **Enable cache replication** check box in the **Cache replication** field.

To manage **batch update** or **PUSH-PULL - PUSH/PULL**, repeat steps 1-4, then, click the **Advanced** button in the Cache replication field. Batch update interval is set under **push frequency**. PUSH-PULL-PUSH/PULL is set through the **runtime mode**.

You can also select which replication domain and initial replicator entry the dynamic cache will utilize (either those managed within the cell or across the cell).

## Cache replication

Data is generated one time and copied or replicated to other servers in the cluster, thus saving execution time and resources. Caching in a cluster has additional concerns. In particular, the same data could be required, and hence, generated in multiple places. Also, the access the resources need to generate the cached data can be restricted, preventing access to the data.

Cache replication addresses these concerns by generating the data one time and copying or replicating it to the other servers in the cluster. It also aids in cache consistency, in that cache entries that are not needed are removed or replaced.

The configuration specific to replication of data can exist as part of the Web container dynamic cache configuration accessible through the administrative console, or on a per cache entry basis through the `cachespec.xml` file. This includes the option to configure cache replication at the Web container level, but disabling it for a specific cache entry.

Cache replication can take on three forms:

- **PUSH** - Send out new entries, both ID and data, and updates to those entries.
- **PULL** - Requests data from other servers in the cluster when that data is not locally present.
- **PUSH/PULL** - Sends out IDs for new entries, then, only request from other servers in the cluster entries for IDs previously broadcast. The dynamic cache always sends out cache entry invalidations.

The dynamic cache provides a batch update option. Specifically, for PUSH or PUSH/PULL, the dynamic cache broadcasts the update asynchronously, based on a



timed interval rather than sending them immediately upon inception. Invalidators are sent immediately. Distribution of invalidations addresses the issue of stale data residing in a cluster.

## Internal messaging configuration settings

Use this page to set advanced configurations for Memory to Memory session replication.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > Cache replication > Enable cache replication** .

The advanced replication settings include fields for choosing the initial replicator entry that connects to the replicator domains. As an alternative, you can specify the IP addresses and ports (of the form address:port) for connection to replicators outside of the cell that the server is administered under. By default, if a replicator is defined on the server you are configuring, that server is the one chosen for cache replication. Select the advanced properties only if you want to deviate from the default setting.

### Internal messaging server

Specifies a domain from which your data will be replicated. Depending on the domain you choose to replicate the data, you can choose any of the replicators defined under that domain. You can use the default domain or choose one from the drop down window.

### Runtime mode

Specifies the global sharing policy for this server.

The following settings are available:

- **Both push and pull** sends the cache ID of newly updated content to other servers in the replication domain. Then, if one of the other servers requests the content, and that server has the ID of the cache entry for the previously updated content, it will retrieve the content from the publishing server. On the other hand, if a request is made for an ID which has not been previously published, the server assumes it does not exist in the cluster and creates a new entry.
- **Push only** sends the cache ID and cache content of new content to all other servers in the replication domain.
- The sharing policy of **Not Shared** results in the cache ID and cache content not being shared with other servers in the replication domain.

The default setting for a non-clustered environment is **Not Shared**. When enabling replication, the default value is **Push only**.

### Push frequency

Specifies the time in seconds to wait before pushing new or modified cache entries to other servers. A value of 0 (zero) means send immediately.

Setting this property to a value greater than 0 (zero) causes a "batch" push of all cache entries that are created or modified during the time period.

|         |                             |
|---------|-----------------------------|
| Default | 0 (equivalent to immediate) |
|---------|-----------------------------|

## Cache replication

Data is generated one time and copied or replicated to other servers in the cluster, thus saving execution time and resources. Caching in a cluster has additional concerns. In particular, the same data could be required, and hence, generated in multiple places. Also, the access the resources need to generate the cached data can be restricted, preventing access to the data.

Cache replication addresses these concerns by generating the data one time and copying or replicating it to the other servers in the cluster. It also aids in cache consistency, in that cache entries that are not needed are removed or replaced.

The configuration specific to replication of data can exist as part of the Web container dynamic cache configuration accessible through the administrative console, or on a per cache entry basis through the `cachespec.xml` file. This includes the option to configure cache replication at the Web container level, but disabling it for a specific cache entry.

Cache replication can take on three forms:

- **PUSH** - Send out new entries, both ID and data, and updates to those entries.
- **PULL** - Requests data from other servers in the cluster when that data is not locally present.
- **PUSH/PULL** - Sends out IDs for new entries, then, only request from other servers in the cluster entries for IDs previously broadcast. The dynamic cache always sends out cache entry invalidations.

The dynamic cache provides a batch update option. Specifically, for PUSH or PUSH/PULL, the dynamic cache broadcasts the update asynchronously, based on a timed interval rather than sending them immediately upon inception. Invalidators are sent immediately. Distribution of invalidations addresses the issue of stale data residing in a cluster.

---

## Enabling the dynamic cache service

In order to use the dynamic cache service, you must first enable it.

### Steps for this task

1. Open the administrative console.
2. Click **Servers** > **Application Servers** in the administrative console navigation tree.
3. Click a server.
4. Click **Dynamic Cache Service** under Additional Properties.
5. Select **Enable service at server startup** in the **Startup state** field.
6. Click **Apply** or **OK**.
7. Restart WebSphere Application Server.

### Results

The dynamic cache service will now cache content for requests that have cache policies configured.

## Dynamic cache service settings

Use this page to configure and manage the dynamic cache service settings.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service**.

### Startup state

Specifies whether the dynamic cache is enabled.

### Cache size

Specifies a positive integer as the value for the maximum number of entries the cache holds.

Enter the cache size value in this field between the range of 100 through 200,000.

### Default priority

Specifies the default priority for cache entries, determining how long an entry stays in a full cache.

|         |   |
|---------|---|
| Default | 1 |
|---------|---|

### Disk offload

Specifies whether disk offload is enabled.

By default, the dynamic cache maintains the number of entries configured in memory. If new entries are created while the cache is full, the priorities configured for each cache entry and a least recently used algorithm, are used to remove entries from the cache. As an alternative, you can configure disk offload, and rather than remove the entries from memory, have them copied onto the file system (the location is configurable) for later.

### Cache replication

Specifies whether cache replication is enabled.

You can also configure advanced cache replication settings.

## Configuring servlet caching

### Before you begin

To enable servlet caching, you must enable the dynamic cache service.

### Steps for this task

1. Open the administrative console.
2. Click **Servers > Application Servers** in the console navigation tree.
3. Click a server.
4. Click **Web Container**.
5. Select the **Enable servlet caching** check box under the Configuration tab.
6. Click **Apply** or **OK**.

### Servlet caching

After a servlet is invoked and generating the output to cache, a cache entry is created containing the output and the side effects of the invocation. For example, these side effects can include calls to other servlets or JavaServer Pages (JSP) files, as well as metadata about the entry, including timeout and entry priority information.

Unique entries are distinguished by an ID string generated from the `HttpServletRequest` object for each invocation of the servlet. You can then base servlet caching on:

- Request parameters and attributes the URI used to invoke the servlet
- Session information
- Other options, including cookies

Since JSP files are compiled by WebSphere Application Server into servlets, the dynamic cache function treats them the same, except in specifically documented situations.

## Configuring cache replication

Cache replication leverages the WebSphere internal replication service that is also leveraged for `HttpSession` memory-to-memory replication for failover purposes. Hence, a replication domain with at least one replicator entry needs to exist in order to replicate the data. The dynamic cache service, in essence, connects to the replicator. See more information in the topic referring to managing internal replication.

To configure cache replication and its features:

### Steps for this task

1. Click **Servers>Manage Application Servers** in the administrative console navigation tree.
2. Click *server*.
3. Click **Dynamic Cache Service**.
4. Click **Enable cache replication** check box in the **Cache replication** field.

To manage **batch update** or **PUSH-PULL - PUSH/PULL**, repeat steps 1-4, then, click the **Advanced** button in the Cache replication field. Batch update interval is set under **push frequency**. PUSH-PULL-PUSH/PULL is set through the **runtime mode**.

You can also select which replication domain and initial replicator entry the dynamic cache will utilize (either those managed within the cell or across the cell).

### Cache replication

Data is generated one time and copied or replicated to other servers in the cluster, thus saving execution time and resources. Caching in a cluster has additional concerns. In particular, the same data could be required, and hence, generated in multiple places. Also, the access the resources need to generate the cached data can be restricted, preventing access to the data.

Cache replication addresses these concerns by generating the data one time and copying or replicating it to the other servers in the cluster. It also aids in cache consistency, in that cache entries that are not needed are removed or replaced.

The configuration specific to replication of data can exist as part of the Web container dynamic cache configuration accessible through the administrative console, or on a per cache entry basis through the `cachespec.xml` file. This includes the option to configure cache replication at the Web container level, but disabling it for a specific cache entry.

Cache replication can take on three forms:

- **PUSH** - Send out new entries, both ID and data, and updates to those entries.
- **PULL** - Requests data from other servers in the cluster when that data is not locally present.
- **PUSH/PULL** - Sends out IDs for new entries, then, only request from other servers in the cluster entries for IDs previously broadcast. The dynamic cache always sends out cache entry invalidations.

The dynamic cache provides a batch update option. Specifically, for **PUSH** or **PUSH/PULL**, the dynamic cache broadcasts the update asynchronously, based on a timed interval rather than sending them immediately upon inception. Invalidators are sent immediately. Distribution of invalidations addresses the issue of stale data residing in a cluster.

## Internal messaging configuration settings

Use this page to set advanced configurations for Memory to Memory session replication.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > Cache replication > Enable cache replication** .

The advanced replication settings include fields for choosing the initial replicator entry that connects to the replicator domains. As an alternative, you can specify the IP addresses and ports (of the form address:port) for connection to replicators outside of the cell that the server is administered under. By default, if a replicator is defined on the server you are configuring, that server is the one chosen for cache replication. Select the advanced properties only if you want to deviate from the default setting.

**Internal messaging server:** Specifies a domain from which your data will be replicated. Depending on the domain you choose to replicate the data, you can choose any of the replicators defined under that domain. You can use the default domain or choose one from the drop down window.

**Runtime mode:** Specifies the global sharing policy for this server.

The following settings are available:

- **Both push and pull** sends the cache ID of newly updated content to other servers in the replication domain. Then, if one of the other servers requests the content, and that server has the ID of the cache entry for the previously updated content, it will retrieve the content from the publishing server. On the other hand, if a request is made for an ID which has not been previously published, the server assumes it does not exist in the cluster and creates a new entry.
- **Push only** sends the cache ID and cache content of new content to all other servers in the replication domain.
- The sharing policy of **Not Shared** results in the cache ID and cache content not being shared with other servers in the replication domain.

The default setting for a non-clustered environment is **Not Shared**. When enabling replication, the default value is **Push only**.

**Push frequency:** Specifies the time in seconds to wait before pushing new or modified cache entries to other servers. A value of 0 (zero) means send immediately.

Setting this property to a value greater than 0 (zero) causes a "batch" push of all cache entries that are created or modified during the time period.

|         |                             |
|---------|-----------------------------|
| Default | 0 (equivalent to immediate) |
|---------|-----------------------------|

## Cache replication

Data is generated one time and copied or replicated to other servers in the cluster, thus saving execution time and resources. Caching in a cluster has additional concerns. In particular, the same data could be required, and hence, generated in multiple places. Also, the access the resources need to generate the cached data can be restricted, preventing access to the data.

Cache replication addresses these concerns by generating the data one time and copying or replicating it to the other servers in the cluster. It also aids in cache consistency, in that cache entries that are not needed are removed or replaced.

The configuration specific to replication of data can exist as part of the Web container dynamic cache configuration accessible through the administrative console, or on a per cache entry basis through the `cachespec.xml` file. This includes the option to configure cache replication at the Web container level, but disabling it for a specific cache entry.

Cache replication can take on three forms:

- **PUSH** - Send out new entries, both ID and data, and updates to those entries.
- **PULL** - Requests data from other servers in the cluster when that data is not locally present.
- **PUSH/PULL** - Sends out IDs for new entries, then, only request from other servers in the cluster entries for IDs previously broadcast. The dynamic cache always sends out cache entry invalidations.

The dynamic cache provides a batch update option. Specifically, for PUSH or PUSH/PULL, the dynamic cache broadcasts the update asynchronously, based on a timed interval rather than sending them immediately upon inception. Invalidators are sent immediately. Distribution of invalidations addresses the issue of stale data residing in a cluster.

## Configuring the dynamic cache disk offload

By default, when the number of cache entries reaches the configured limit for a given WebSphere server, eviction of cache entries occurs, allowing new entries to enter the cache service. The dynamic cache includes an alternative feature named disk offload, that copies the evicted cache entries to disk for potential future access.

To configure disk offload:

### Steps for this task

1. Open the administrative server.
2. Click **Server > Application Server** in the administrative console navigation tree.
3. Click *server*.
4. Click **Dynamic Cache Service**.
5. Click the **Enable disk offload** check box in the **Disk offload** field.  
You can also set the disk offload location in this field.
6. Click **Apply** or **OK**.

## Application servers must have different disk offload locations

When you have two or more application servers with servlet caching enabled and the application servers specify the same disk offload location for their caches through the dynamic cache service, the following exceptions might occur:

```
java.lang.NullPointerException
    at com.ibm.ws.cache.CacheOnDisk.readTemplate(CacheOnDisk.java:686)
    at com.ibm.ws.cache.Cache.internalInvalidateByTemplate(Cache.java:828)
```

or:

```
java.lang.NullPointerException
    at com.ibm.ws.cache.CacheOnDisk.readCacheEntry(CacheOnDisk.java:600)
    at com.ibm.ws.cache.Cache.getCacheEntry(Cache.java:341)
```

If one server is run as root and the other servers are run as nonroot, this problem could occur. For example, if server1 runs as root and server2 runs as wasuser or wasgroup, the cache files in the disk offload location might be created with root permissions. This situation causes the applications running on the nonroot servers to crash when they try to read or write to the cache.

The disk offload location must be unique for servers defined on the same node. If you have multiple servers defined on the same node, make sure the disk offload location is different for each server as defined on the **Dynamic Cache Service** panel, **Offload location** field.

## Configuring Edge Side Include caching

Edge Side Include (ESI) is configured through the plugin-cfg.xml file.

The Web server plug-in contains a built-in ESI processor. The ESI processor has the ability to cache whole pages, as well as fragments, providing a higher cache hit ratio. The cache implemented by the ESI processor is an in-memory cache, not a disk cache, therefore, the cache entries are not saved when the Web server is restarted.

The basic operation of the ESI processor is as follows: When a request is received by the Web server plug-in, it is sent to the ESI processor, unless the ESI processor is disabled. It is enabled by default. If a cache miss occurs, a Surrogate-Capabilities header is added to the request and the request is forwarded to the WebSphere Application Server. If the dynamic servlet cache is enabled in the application server, and the response is edge cacheable, the application server returns a Surrogate-Control header in response to the WebSphere Application Server plug-in.

The value of the Surrogate-Control response header contains the list of rules which are used by the ESI processor in order to generate the cache ID. The response is then stored in the ESI cache, using the cache ID as the key. For each ESI include tag in the body of the response, a new request is processed such that each nested include results in either a cache hit or another request forwarded to the application server. When all nested includes have been processed, the page is assembled and returned to the client.

The ESI processor is configurable through the WebSphere Web server plug-in configuration file plugin-cfg.xml. The following is an example of the beginning of this file, which illustrates the ESI configuration options.

```
<?xml version="1.0"?>
<Config>
  <Property Name="esiEnable" Value="true"/>
  <Property Name="esiMaxCacheSize" Value="1024"/>
  <Property Name="esiInvalidationMonitor" Value="false"/>
</Config>
```

The first option, `esiEnable`, can be used to disable the ESI processor by setting the value to false. ESI is enabled by default. If ESI is disabled, then the other ESI options are ignored.

The second option, `esiMaxCacheSize`, is the maximum size of the cache in 1K byte units. The default maximum size of the cache is 1 megabyte. If the cache is full, the first entry to be evicted from the cache is the entry that is closest to expiration.

The third option, `esiInvalidationMonitor`, specifies whether or not the ESI processor should receive invalidations from the application server. ESI works well when the Web servers following a threading model is used, and only one process is started. When multiple processes are started, each process caches the responses independently and the cache is not shared. This could lead to a situation where, the system's memory is fully used up by ESI processor. There are three methods by which entries are removed from the ESI cache: first, an entry's expiration timeout could fire; second, an entry may be purged to make room for newer entries; or third, the application server could send an explicit invalidation for a group of entries. In order for the third mechanism to be enabled, the `esiInvalidationMonitor` property must be set to true and the `DynaCacheEsi` application must be installed on the application server. The `DynaCacheEsi` application is located in the `installableApps` directory and is named `DynaCacheEsi.ear`. If the `ESIInvalidationMonitor` property is set to true but the `DynaCacheEsi` application is not installed, then errors will occur in the webserver plugin and the request will fail.

The ESI processor's cache can be monitored through the `CacheMonitor` application. In order for ESI processor's cache to be visible in the `CacheMonitor`, the `DynaCacheEsi` application must be installed as described above and the `ESIInvalidationMonitor` property must be set to true in the `plugin-cfg.xml` file.

When WebSphere Application Server is used to serve static data, such as images and HTML on the application server, the URLs are also cached in the ESI processor. This data has a default timeout of 300 seconds. You can change the timeout value by adding the property `com.ibm.servlet.file.esi.timeOut` to your JVM's command line parameters. The following example shows how to set a one minute timeout on static data cached in the plug-in:

```
-Dcom.ibm.servlet.file.esi.timeOut=60
```

For more information about the `plugin-cfg.xml` file see the topic "Improving performance with the dynamic cache service".

**5.0.1** For information about configuring alternate URL, see the topic [Configuring alternate URL](#).

### Configuring alternate URL

Alternate URL is a method for edge caching JavaServer Pages (JSP) files and servlet responses that you can not request externally. Dynamic cache provides support to recognize the presence of an Edge Side Include (ESI) processor and to generate ESI include tags and appropriate cache policies for edge cacheable fragments. However, for a fragment to be edge cacheable, you must be able to



externally request it from the application server. In other words, if a user types the URL in their browser with the appropriate parameters and cookies for the fragment, WebSphere Application Server must return the content for that fragment.

One of the standard J2EE programming architectures is the model-view-controller (MVC) architecture, where a call to a controller servlet might include one or more child JSP files to construct the view. When using the MVC programming model, the child JSP files are edge cacheable only if you can request these JSP files externally, which is not usually the case. For example, if a child JSP file uses one or more request attributes that are determined and set by the controller servlet, you cannot cache that JSP file on the edge. You can use alternate URL support to overcome this limitation by providing an alternate controller servlet URL used to invoke the JSP file.

The alternate URL for a JSP file or a servlet is set in the `cachespec.xml` file as a property with the name `alternate_url`. You can set the alternate URL either on a per cache-entry basis or on a per cache-id basis. It is valid only if the `EdgeCacheable` property is also set for that entry. If the `EdgeCacheable` property is not set, the `alternate_url` property is ignored. The following is a sample cache policy using the `alternate_url` property:

```
<cache-entry>
  <class>servlet</class>
  <name>/AltUrlTest2.jsp</name>
  <property name="EdgeCacheable">>true</property>
  <property name="alternate_url">/alturlcontroller2</property>
  <cache-id>
    <timeout>600</timeout>
    <priority>2</priority>
  </cache-id>
</cache-entry>
```

### What to do next

For more information on the `cachespec.xml` file, see `Cachespec.xml` file.

## Configuring external cache groups

The dynamic cache can control caches outside of the application server, such as IBM Edge Server, an IBM HTTP Server for distributed platforms' Fast Response Cache Accelerator (FRCA) cache, and a WebSphere HTTP Server for distributed platforms plug-in ESI Fragment Processor. When external cache groups are defined, the dynamic cache matches externally cacheable cache entries with those groups, and pushes cache entries and invalidations out to those groups. This allows WebSphere Application Server to manage dynamic content beyond the application server. The content can then be served from the external cache, instead of the application server, improving savings in performance.

### Steps for this task

1. Open the administrative console.
2. Enable the dynamic cache.
  - a. Click **Servers > Application Servers** in the administrative console navigation tree.
  - b. Click a *server*.
  - c. Click **Dynamic Cache Service**.
  - d. Select the check box in the **Startup state** field to enable the dynamic cache.

3. Define the external cache group in which WebSphere Application Server should control.
  - a. Click **External Caching Groups** from the Dynamic Cache administrative console page.
  - b. Click **New** or choose an external cache group from the list.
4. Configure cache group members.
  - a. Click **External cache groups** from the Dynamic Cache administrative console page. Then click **New** or choose an external cache group from the list.
  - b. Click **External cache group members > New** or choose an external cache group member from the list.
  - c. Type the configuration string in the **Address** field.
  - d. Type the adapter bean name in the **Adapter Bean Name** field.
  - e. **Save** the configuration.
  - f. Click **Apply** or **OK**.

### External cache group collection

Use this page to define sets of external caches controlled by WebSphere Application Server on Web servers, such as IBM Edge Server and IBM HTTP Server.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache Groups**.

**Name:** Specifies the external cache group name.

The external cache group name needs to match the externalcache property as defined in the servlet or JSP cachespec.xml file.

When external caching is enabled, the cache matches pages with its URIs and pushes matching pages to the external cache. The entries can then be served from the external cache, instead of the application server.

**Type:** Specifies the external cache group type.

### External cache group settings

Use this page to configure sets of external caches controlled by WebSphere Application Server on Web servers, such as IBM Edge Server and IBM HTTP Server.

To view this administrative console page, click **Servers > Application Server > server > Dynamic Cache Service > External Cache groups > external\_cache\_group**.

**Name:** Specifies the external cache group name.

The external cache group name needs to match the externalcache property as defined in the servlet or JavaServer Pages (JSP) cachespec.xml file.

When external caching is enabled, the cache matches pages with its URIs and pushes matching pages to the external cache. The entries can then be served from the external cache, instead of the application server. This ability creates a significant savings in performance.

**Type:** Specifies the external cache group type.

## External cache group member collection

Use this page to define specific caches that are members of a cache group.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache groups > external\_cache\_group > External cache group members**.

**Address:** Specifies a configuration string used by external cache adapter bean to connect to the external cache.

**AdapterBeanName:** Specifies the adapter bean name.

Example adapter bean names supported in WebSphere Application Server are:

### AFPA

AdapterBeanName: com.ibm.ws.cache.servlet.Afpa

Address: Port on which afpa listens

&nbsp;

### ESI

AdapterBeanName: com.ibm.websphere.servlet.cache.ESIInvalidatorServlet

Address: local host

&nbsp;

### WTE

AdapterBeanName: com.ibm.websphere.edge.dynacache.WteAdapter

Address: hostname:port (hostname and port on which WTE is listening)

## External cache group member settings

Use this page to configure specific caches that are members of a cache group.

To view this administrative console page, click **Servers > Application Servers > server > Dynamic Cache Service > External Cache groups > External\_cache\_group > External cache group members > External\_cache\_group\_members**.

**Address:** Specifies a configuration string used by external cache adapter bean to connect to the external cache.

**AdapterBeanName:** Specifies the adapter bean name.

### AFPA

AdapterBeanName: com.ibm.ws.cache.servlet.Afpa

Address: Port on which afpa listens

&nbsp;

### ESI

AdapterBeanName: com.ibm.websphere.servlet.cache.ESIInvalidatorServlet

Address: local host

### WTE

AdapterBeanName: com.ibm.websphere.edge.dynacache.WteAdapter

Address: hostname:port (hostname and port on which WTE is listening)

## Configuring high-speed external caching through the Web server

IBM HTTP Server for Windows NT and Windows 2000 operating systems contains a high-speed cache referred to as the *Fast Response Cache Accelerator*, or *cache accelerator*.

The Fast Response Cache Accelerator is available on Windows NT and Windows 2000 operating systems and AIX platforms. However, support to cache dynamic content is only available on Windows NT and Windows 2000 operating systems.

You can enable cache accelerator to cache static and dynamic content. To enable cache accelerator for caching static content, add the following directives to the `http.conf` configuration file, in the IBM HTTP Server conf directory:

- `AfpaEnable`
- `AfpaCache on`
- `AfpaLogFile "install_root\IBMHttpServer\logs\afpalog" V-ECLF`

To enable Cache Accelerator for caching dynamic content, such as servlets and Java Server Pages (JSP) files, configure the WebSphere Application Server and the IBM HTTP Server for distributed platforms:

#### Steps for this task

1. Configure WebSphere Application Server to enable Fast Cache Response Accelerator:
  - a. Turn on servlet caching for each application server that uses the cache accelerator.
  - b. Configure an external cache group on the application server:
    - Click **Servers > Application Servers > server1**.
    - Click **Dynamic Cache Service** in the **Additional Properties** window.
    - Click **External Cache Groups** in the **Additional Properties** window.
    - Click **New** on the External cache group administrative console page to define an external cache group named `afpa` for each application server that uses the cache accelerator.
    - Type `afpa` in the External cache group field.
    - Click **Apply**.
    - Add a member to the group with an adapter bean name of `com.ibm.ws.cache.servlet.Afpa`:

Click **Afpa > External cache group members**. Click **New** on the External cache group members administrative console page. Type `com.ibm.ws.cache.servlet.Afpa` in the AdapterBean name field. Enter an unused port number in the Address field.

- c. Add a cache policy in the `cachespec.xml` file for the servlet or JSP file you want to cache. Add the following property to the cache policy:

```
<property name="ExternalCache">afpa</property>
```

It is important to follow all the steps for every application server in the cluster.

2. Enable cache accelerator on the IBM HTTP Server for distributed platforms:
  - a. Add the following directives to the end of the `httpd.conf` file:
    - `AfpaEnable`
    - `AfpaCache on`
    - `AfpaLogFile "install_root\IBMHttpServer\logs\afpalog" V-ECLF`
    - `LoadModule afpaplugin_module install_root/bin/afpaplugin.dll`
    - `AfpaPluginHost WAS_Hostname:port`, where `WAS_Hostname` is the host name of the application server and `port` is the port you specified in the Address field while configuring the external cache group member

The LoadModule directive loads the IBM HTTP Server plug-in that connects the Fast Response Cache Accelerator to the WebSphere Application Server fragment cache.

If multiple IBM HTTP Servers are routing requests to a single application server, add the directives above to the http.conf file of each of these IBM HTTP Servers for distributed platforms.

If one IBM HTTP Server is routing requests to a cluster of application servers, add the AfpapluginHost WAS\_Hostname:port directive to the http.conf file for each application server in the cluster. For example, if there are three application servers in the cluster, add the following directives to the http.conf file:

- LoadModule afpaplugin\_module *install\_root*/bin/afpaplugin.dll
- AfpapluginHost WAS1\_Hostname:port1
- AfpapluginHost WAS2\_Hostname:port2
- AfpapluginHost WAS3\_Hostname:port3

### **Configuring Fast Response Cache Accelerator cache size through a distributed platforms Web server**

In the default IBM HTTP Server for distributed platforms configuration, the maximum Fast Cache Accelerator dynamic cache size is calculated as 1/8 of physical pin-able memory. On a machine with 384MG of RAM, it allows a maximum of approximately 50MG for the Fast Cache Accelerator dynamic cache. When this limit is reached, the Cache Accelerator then deletes older entries to cache new ones.

Follow these steps to configure the Cache Accelerator:

#### Steps for this task

1. Using the IBM HTTP Server for distributed platforms' AfpadynaCacheMax directive, tune the maximum allowed cache size:
  - a. Place the directive in the global server configuration scope, along with the other default Fast Cache Accelerator directives.
  - b. Enable Fast Cache Accelerator.

To enable the Fast Cache Accelerator, update the following directives in this IBM HTTP Server's http.conf file:

```
AfpaEnable
AfpaCache on
Afpalogs "c:/Program Files/IBM HTTP Server/logs/afpalogs" V-ECLF
AfpadynaCacheMax 10
```

These above settings limit the dynamic cache size to 10MG. If you use these directives to increase cache size, do not make the cache so large that all the physical memory is consumed. Determine how much memory is available when all applications are running, by using the Windows Task Manager.

Assign no more than 50% of available physical memory to the dynamic cache. Specifying too large a cache not only decreases the performance of other applications, but also puts you at a risk for completely running out of memory.

The default configuration does not include the AfpadynaCacheMax directive where the cache size is automatically calculated as 1/8 of physical memory.

---

## Displaying cache information

The dynamic cache monitor is an installable Web application that displays simple cache statistics, cache entries, and cache policy information. Follow the steps below to install the cache monitor application.

### Steps for this task

1. Install the cache monitor application from the *install\_root/installableApps* directory.  
The application is named CacheMonitor.ear.
2. Access the Web application using a Web browser and the URL `http://your_host_name:your_port_number/cachemonitor`, where *your port number* is the port associated with the host on which you installed the cache monitor application.

---

## Configuring cacheable objects with the cachespec.xml file

### Before you begin

Define cacheable objects inside the `cachespec.xml`, found inside the Web module WEB-INF or enterprise bean META-INF directory.

You can place a global `cachespec.xml` in the application server properties directory, but the recommended method is to place the cache configuration file with the deployment module. The root element of the `cachespec.xml` file is `<cache>`, which contains `<cache-entry>` elements.

Within a `<cache-entry>...</cache-entry>` element are parameters that allow you to complete the following tasks to enable the dynamic cache with the `cachespec.xml` file:

### Steps for this task

1. Develop a `cachespec.xml` file.
  - a. Create a caching configuration file.  
In the `<install_root>/properties` directory, locate the `cachespec.sample.xml` file.
  - b. Copy the `cachespec.sample.xml` file to `cachespec.xml` in Web module WEB-INF or enterprise bean META-INF directory.
2. Define the cache-entry elements necessary to identify the cacheable objects. See the topic "Cachespec.xml file" for a list of elements.
3. Develop cache-ID rules.

To cache an object, WebSphere Application Server must know how to generate unique IDs for different invocations of that object. The `<cache-id>` element performs that task. Each cache entry can have multiple cache-ID rules that execute in order until either a rule returns non-empty cache-ID or no more rules remain to execute. If none of the cache-ID generation rules produce a valid cache ID, then the object is not cached. Develop these IDs in one of two ways:

- Use the `<component>` element defined in the cache policy of a cache entry (recommended)
- Write custom Java code to build the ID from input variables and system state

To configure the cache entry to use the IdGenerator, specify your IdGenerator in the XML file by using the `<idgenerator>` tag, for example:

```
<cache-entry>
  <class>servlet</class>
  <name>/servlet/CommandProcessor</name>
<cache-id>
  <idgenerator>com.mycompany.SampleIdGeneratorImpl</idgenerator>
  <timeout>60</timeout>
</cache-id>
</cache-entry>
```

You can also use the Application Assembly Tool (AAT) to define the IdGenerator class in the cache policy's **Advanced** tab.

#### 4. Specifying dependency ID rules.

Use dependency ID elements to specify additional cache group identifiers that associate multiple cache entries to the same group identifier.

The dependency ID is generated by concatenating the dependency ID base string with the values returned by its component elements. If a required component returns a null value, then the entire dependency ID does not generate and is not used. You can validate the dependency IDs explicitly through the WebSphere Dynamic Cache API, or use another cache-entry `<invalidation>` element. Multiple dependency ID rules can exist per cache-entry. All dependency ID rules separately execute. See the topic "Cachespec.xml file" for a list of `<component>` elements.

#### 5. Invalidate other cache entries as a side effect of this object execution, if relevant.

You can define invalidation rules in exactly the same manner as dependency IDs. However, the IDs that generate by invalidation rules are used to invalidate cache entries that have those same dependency IDs.

The invalidation ID is generated by concatenating the invalidation ID base string with the values returned by its component element. If a required component returns a null value, then the entire invalidation ID is not generated and no invalidation occurs. Multiple invalidation rules can exist per cache-entry. All invalidation rules separately execute.

#### 6. Verify the cacheable page.

### What to do next

Typically you declare several `<cache-entry>...</cache-entry>` elements inside a `cachespec.xml` file.

The dynamic cache responds to changes in this file. When new versions of the `cachespec.xml` are detected, the old policies are replaced. Objects cached through the old policy file are not automatically invalidated from the cache; they are either reused with the new policy or eliminated from the cache through its replacement algorithm.

For each of the three IDs (cache, dependency, invalidation) generated by cache entries, a `<cache-entry>` can contain multiple elements. The dynamic cache will execute the `<cache-id>` rules in order, and the first one that successfully generates an ID will be used to cache that output. If the object is to be cached, each one of the `<dependency-id>` elements will be executed to build a set of dependency IDs for that cache entry. Finally, each of the `<invalidation>` elements will be executed, building a list of IDs that the dynamic cache will invalidate, whether or not this object is cached.

## Verifying the cacheable page

Verify the cacheable page by following these steps:

### Steps for this task

1. View the snoop servlet in the default application by accessing the URI: /snoop
2. Invoke and reload the URI several times using a different Web browser or using different parameters.

This action returns the same output for the snoop servlet. The snoop servlet is now operating incorrectly, because it displays the request information from its first invocation rather than from the current request.

3. Inspect the entry in the cache with the dynamic cache monitor.

## Cachespec.xml file

The cache parses the cachespec.xml file on startup, and extracts from each <cache-entry> element a set of configuration parameters. Every time a new servlet or other cacheable object initializes, the cache attempts to match each of the different cache-entry elements, to find the configuration information for that object. Different cacheable objects have different <class> elements. You can define the specific object a cache policy refers to using the <name> element.

### Location

The cachespec.xml file is found inside the WEB-INF directory of a Web module.

You can place a global cachespec.xml file in the application server properties directory, but the recommended method is to place the cache configuration file with the deployment module. (To place the cache configuration file with the deployment module, use the Application Assembly Tool (AAT) to define the cacheable objects.

The root element of the cachespec.xml file is *cache*, which contains *cache-entry* elements.

The cachespec.dtd file is available in the application server properties directory.

### Usage notes

Each cache entry must specify certain basic information that the dynamic cache uses to process that entry. This section explains the function of each cache entry element of the cachespec.xml file including:

- class
- name
- sharing-policy
- property
- cache-id

### class

```
<class>command | servlet | webservice</class>
```

This element is required and governs how the application server interprets the remaining cache policy definition. The value *servlet* refers to servlets and JavaServer Pages (JSP) files deployed in the WebSphere Application Server servlet



engine. The object class extends the servlet with special component types for Web services requests. Finally, the value command refers to classes using the WebSphere command programming model. The following examples illustrate the class element:

```
<class>command</class>
<class>servlet</class>
<class>webservice</class>
```

### name

```
<name>name</name>
```

where *name* is the fully qualified class name of the command, servlet, or object.

There are two ways to use <name> to specify a cacheable object:

- For commands and objects, this required element must include the package name, if any, and class name, including a trailing .class, of the configured object.
- For servlets and JSP files, if the cachespec.xml file is in the WebSphere Application Server properties directory, this required element must include the full URI of the JSP file or servlet to cache. For servlets and JSP files, if the cachespec.xml file is in the Web application, this required element can be relative to the specific Web application context root.

**Note:** The preferred location of the cachespec.xml file is in the Web application, not the properties directory.

You can specify multiple <name> elements within a <cache-entry> if you have different mappings that refer to the same servlet.

The following examples illustrate the name element:

```
<name>com.mycompany.MyCommand.class</name>
<name>default_host:/servlet/snoop</name>
<name>com.mycompany.beans.MyJavaBean</name>
<name>mywebapp/myjsp.jsp</name>
```

### sharing-policy

```
<sharing-policy> not-shared | shared-push | shared-pull </sharing-policy>
```

When working within a cluster with a distributed cache, these values determine the sharing characteristics of entries created from this object. If this element is not present, a not-shared value is assumed. Also, in non-distributed environments, not-shared is the only valid value. This property does not affect distribution to Edge servers through the Edge fragment caching property.

| Value      | Description   |
|------------|---|
| not-shared | Cache entries for this object are not shared among different application servers. These entries can contain non-serializable data. For example, a cached servlet can place non-serializable objects into the request attributes, if the <class> type supports it. |

|                  |  |
|------------------|--|
| shared-push      | Cache entries for this object are automatically distributed to the dynamic caches in other application servers or cooperating Java virtual machines (JVMs). Each cache has a copy of the entry at the time it is created. These entries cannot store non-serializable data.  |
| shared-pull      | Cache entries for this object are shared between application servers on demand. If an application server gets a cache miss for this object, it queries the cooperating application servers to see if they have the object. If no application server has a cached copy of the object, the original application server executes the request and generates the object. These entries cannot store non-serializable data.  |
| shared push-pull | Cache entries for this object are shared between application servers on demand. When an application server generates a cache entry, it broadcasts the cache ID of the created entry to all cooperating application servers. Each server then knows whether an entry exists for any given cache ID. On a given request for that entry, the application server knows whether to generate the entry or pull it from somewhere else. These entries cannot store non-serializable data. |

The following example shows a sharing policy:

```
<sharing-policy>not-shared</sharing-policy>
```

### property

```
<property name="key">value</property>
```

where *key* is the name of the property defined for this cache entry element, and *value* is the corresponding value.

You can set optional properties on a cacheable object, such as a description of the configured servlet. The class determines valid properties of the cache entry. At this time, the following properties are defined:

| Property        | Valid classes | Value  |
|-----------------|---------------|--|
| ApplicationName | All           | Overrides the J2EENAME application ID so that multiple applications can share a common cache ID namespace. |

|                      |                        |  |
|----------------------|------------------------|--|
| EdgeCacheable        | Servlet                | True or false. Default is false. If the property is true, then the given servlet or JSP file is externally requested from an Edge Server. Whether or not the servlet or JSP file is cacheable, depends on the rest of the cache specification.   |
| ExternalCache        | Servlet                | Specifies the external cache name. The external cache name needs to match the external cache group name.   |
| consume-subfragments | Servlet or Web service | True or false. Default is false. When a servlet is cached, only the content of that servlet is stored, and includes placeholders for any other fragments to which it includes or forwards. Consume-subfragments (CSF) tells the cache not to stop saving content when it includes a child servlet. The parent entry, the one marked CSF, includes all the content from all fragments in its cache entry, resulting in one big cache entry that has no includes or forwards, but the content from the whole tree of entries. This can save a significant amount of application server processing, but is typically only useful when the external HTTP request contains all the information needed to determine the entire tree of included fragments. |

### cache-id

To cache an object, the application server must know how to generate a unique ID for different invocations of that object. These IDs are built either from user-written custom Java code or from rules defined in the cache policy of each cache entry. Each cache entry can have multiple cache ID rules that are executed in order until either:

- A rule returns a non-empty cache ID, or
- No more rules are left to execute.

If none of the cache ID generation rules produce a valid cache ID, the object is not cached.

Each cache-id element defines a rule for caching an object and is composed of the sub-elements component, timeout, priority, and property. The following example illustrates a cache-id:

<cache-id>component\*| timeout? | priority? | property\* </cache-id>

### component sub-element

Use the component sub-element to generate a portion of the cache ID. Each component sub-element consists of the attributes `id`, `type`, and `ignore-value`, and the elements `method`, `field`, `required`, `value`, and `not-value`.

- Use the `id` attribute to identify the component.
- Use the `type` attribute to identify the type of component. The following table lists the values for the type.

| Type                      | Valid classes           | Meaning   |
|---------------------------|-------------------------|---|
| method                    | command                 | Calls the indicated method on the command or object   |
| field                     | command                 | Retrieves the named field in the command or object  |
| parameter                 | servlet                 | Retrieves the named parameter value from the request object   |
| parameter-list            | servlet                 | Retrieves a list of values for the named parameter  |
| session                   | servlet                 | Retrieves the named value from the HttpSession  |
| cookie                    | servlet                 | Retrieves the named cookie value  |
| attribute                 | servlet                 | Retrieves the named request attribute   |
| header                    | servlet and Web service | Retrieves the named request header  |
| pathInfo                  | servlet                 | Retrieves the pathInfo from the request   |
| servletpath               | servlet                 | Retrieves the servlet path  |
| locale                    | servlet                 | Retrieves the request locale  |
| SOAPEnvelope              | Web service             | Retrieves the SOAPEnvelope from a Web services request. An ID attribute of Hash uses a Hash of the SOAPEnvelope, while Literal uses the SOAPEnvelope as received. |
| SOAPAction                | Web service             | Retrieves the SOAPAction header, (if available), for a Web services request.  |
| serviceOperation          | Web service             | Retrieves the service operation for a Web services request  |
| serviceOperationParameter | Web service             | Retrieves the specified parameter from a Web services request   |

- Use the `ignore-value` attribute to specify whether or not to use the value returned by this component in cache ID formation. This is an optional attribute with a default value of `false`. If the value is `true`, only the ID of the component is used when creating a cache ID, or no output is used when creating a dependency or invalidation ID.
- Use the **method** element to call a void method on a returned object. You can infinitely nest `method` and `field` objects in any combination. The `method` must be public and is not valid for edge-cacheable components. For example:

```
<component id="getUser" type="method"><method>getUserInfo
<method>getName</method></method></component>
```

This method is equivalent to `getUser().getUserInfo().getName()`

- Use the **field** element to access a field in a returned object. You can infinitely nest method and field objects in any combination. The field must be public. Not valid for edge-cacheable components. For example:

```
<component id="getUser" type="method"><method>getUserInfo
<field>name</field></method></component>
```

This method is equivalent to `getUser().getUserInfo().name`

- Use the **required** element to specify whether or not this component must return a non-null value for this cache ID for it to represent a valid cache. If set to `true`, this component must return a non-null value for this cache ID to represent a valid cache ID. If set to `false`, the default, a non-null value is used in the formation of the cache ID and a null value means that this component is not used at all in the ID formation. For example:

```
<required>true</required>
```

- Use the **value** element to specify values that must match to use this component in cache ID formation. For example:

```
<component id="getUser" type="method"><value>blue</value>
<value>red</value> </component>
```

- Use the **not-value** element to specify values that must not match to use this component in cache ID formation. This method is similar to `<value>`, but instead prescribes the defined values from caching. You can use multiple `<not-value>` elements when there is more than one invalid value. For example:

```
<component id="getUser" type="method">
<required>true</required>
<not-value>blue</not-value>
<not-value>red</not-value></component>
```

The component element can have either a method or a field element, or a value or a not-value element. The method and field elements apply only to commands. The following example illustrates the attributes of a component element:

```
<component id="isValid" type="method" ignore-value="true"><component>
```

### timeout sub-element

The timeout sub-element is used to specify a time-to-live (TTL) value for the cache entry. For example,

```
<timeout>value</timeout>
```

where *value* is the amount of time, in seconds, to keep the cache entry. If 0, or a negative value is specified, the cache entry is kept indefinitely.

### priority sub-element

Use the priority sub-element to specify the priority of a cache entry in a cache. The priority weighting is used by the least recently used (LRU) algorithm, of the cache to decide which entries to remove from the cache if the cache runs out of storage space. For example,

```
<priority>value</priority>
```

where *value* is a positive integer between 1 and 255 inclusive.

### property sub-element

Use the property sub-element to specify generic properties for the cache entry. For example,

```
<property name="key">value</property>
```

where *key* is the name of the property to define, and *value* is the corresponding value.

For example:

```
<property name="description">The Snoop Servlet</property>
```

| Property                        | Valid classes | Meaning   |
|---------------------------------|---------------|---|
| sharing-policy/timeout/priority | All           | Overrides the settings for the containing cache entry when the request matches this cache ID. |
| EdgeCacheable                   | servlet       | Overrides the settings for the containing cache entry when the request matches this cache ID. |

### idgenerator and metadatagenerator elements

Use the idgenerator element to specify the class name loaded for the generation of the cache ID. The IdGenerator must implement the `com.ibm.websphere.servlet.cache.IdGenerator` interface. The IdGenerator must build and set cache IDs, group IDs and invalidation IDs. An example of the idgenerator element follows:

```
<idgenerator> classname classname </idgenerator>
```

(where *classname*= Fully qualified name of the class to use)

Use the metadatagenerator element to specify the class name loaded for the metadata generation cache ID. The MetadataGenerator class must implement the `com.ibm.websphere.servlet.cache.MetaDataGenerator` interface. The MetadataGenerator defines properties like timeout, external caching or generic properties. An example of the metadatagenerator element follows:

```
<metadatagenerator> classname classname </metadatagenerator>
```

(where *classname*= Fully qualified name of the class to use)

---

## Configuring command caching

Cacheable commands are stored in the cache for re-use with a similar mechanism for servlets and Java Server Pages (JSP) files. However, in this case, the unique cache IDs are generated based on methods and fields present in the command as input parameters. For example, a **GetStockQuote** command can have a symbol as its input parameter.

A unique cache ID can generate from the name of the command, plus the value of the symbol.

To use command caching you must:

### Steps for this task

#### 1. Create a command.

##### a. Define an interface.

The Command interface specifies the most basic aspects of a command.

You must define the interface that extends one or more of the interfaces in the command package. The command package consists of three interfaces:

- TargetableCommand
- CompensableCommand
- CacheableCommand

In practice, most commands implement the TargetableCommand interface, which allows the command to execute remotely. The code structure of a command interface for a targetable command follows:

```
...
import com.ibm.websphere.command.*;
public interface MyCommand extends TargetableCommand {
    // Declare application methods here
}
```

##### a. Provide an implementation class for the interface.

Write an interface that extends the CacheableCommandImpl class and implements your command interface. This class contains the code for the methods in your interface, the methods inherited from extended interfaces like the CacheableCommand interface, and the required or abstract methods in the CacheableCommandImpl class.

You can also override the default implementations of other methods provided in the CacheableCommandImpl class.

## Command class

To write a command interface, extend one or more of the three interfaces included in the command package. The base interface for all commands is the Command interface. This interface provides only the client-side interface for generic commands and declares three basic methods:

- **isReadyToCallExecute.** This method is called on the client side before the command passes to the server for execution.
- **execute.** This method passes the command to the target and returns any data.
- **reset.** This method reverts any output properties to the values they had before the execute method was called so that you can reuse the object.

The implementation class for your interface must contain implementations for the isReadyToCallExecute and reset methods.

## CacheableCommandImpl class

Commands are implemented by extending the class CacheableCommandImpl, which implements the CacheableCommand interface.

The CacheableCommandImpl class is an abstract class that provides implementations for some of the methods in the CacheableCommand interface, for example, setting return values. This class declares additional methods that the application must implement, for example, how to execute the command.

The code structure of an implementation class for the CacheableCommand interface follows:

```

...
import com.ibm.websphere.command.*;
public class MyCommandImpl extends CacheableCommandImpl
implements MyCommand {
// Set instance variables here      ...
// Implement methods in the MyCommand interface      ...
// Implement abstract methods in the CacheableCommandImpl class
...
}

```

## Example: Caching a command object

This example of command caching is a simple stock quote command.

### Examples

The following is a stock quote command bean. It accepts a ticker as an input parameter and produces a price as its output parameter.

```

public class QuoteCommand extends CacheableCommandImpl
{
    private String ticker;
    private double price;
    // called to validate that command input parameters have been set
    public boolean isReadyToCallExecute() {
        return (ticker!=null);
    }
    // called by a cache-hit to copy output properties to this object
    public void setOutputProperties(TargetableCommand fromCommand) {
        QuoteCommand f = (QuoteCommand)fromCommand;
        this.price = f.price;
    }

    // business logic method called when the stock price must be retrieved
    public void performExecute()throws Exception {...}

    //input parameters for the command
    public void setTicker(String ticker) { this.ticker=ticker;}
    public String getTicker() { return ticker;}

    //output parameters for the command
    public double getPrice() { return price;};
}

```

### Examples

To cache the above command object using the stock ticker as the cache key and using a 60 second time-to-live, use the following cache policy:

```

<cache>
  <cache-entry>
    <class>command</class>
    <sharing-policy>not-shared</sharing-policy>
    <name>QuoteCommand</name>
    <cache-id>
      <component type="method" id="getTicker">
        <required>true</required>
      </component>
      <priority>3</priority>
      <timeout>60</timeout>
    </cache-id>
  </cache-entry>
</cache>

```



---

## Using the DistributedMap interface for the dynamic cache

The DistributedMap interface is a simple interface for the dynamic cache. Using the DistributedMap interface, J2EE applications and system components can cache and share Java objects by storing a reference to the object in the cache. The default WebSphere dynamic cache instance is created if dynamic cache service is enabled in the administrative console. This default instance is bound into the global JNDI namespace using the name `services/cache/distributedmap`. For information on how to enable the dynamic cache service globally, see [Enabling the dynamic cache service](#).

Multiple instances of the DistributedMap interface on the same Java virtual machine (JVM) enables applications to separately configure cache instances as needed. Each instance of DistributedMap has its own properties:

- **Cache JNDI name.** The JNDI name is used for calling a particular DistributedMap instance.
- **Cache size (Property name: `cacheSize`, Default: 2000).** Dynamic cache maintains the maximum number of entries in memory.
- **Enable disk offload (Property name: `enableDiskOffload`, Default: `false`).**
  - Disk offload is disabled - If a new entry is created while the cache is full, the priorities configured for each entry, along with at least recently used algorithm are used to remove the entry from the cache in memory.
  - Disk offload is enabled - If a new entry is created while the cache is full, the priorities configured for each entry, along with a recently used algorithm are used to remove the entry from the cache in memory. The entry that is removed is copied to the local file system, specified by the disk offload location.
- **Disk offload location (Property name: `diskOffloadLocation`).** If disk offload location is not specified, the default location, `$install_root/temp/node/servername/_dynacache/cacheJNDIname` will be used. If disk offload location is specified, the node, server name, and cache instance name will be appended.  
For example, `$install_root/diskoffload` generates the location as `$install_root/diskoffload/node/servername/cacheJNDIname`. This value is ignored if `enableDiskOffload` is `false`.
- **Flush to disk default (Property name: `flushToDiskOnStop`, Default: `false`).** Set this value to `true` to save in-memory cached objects to disk when the server is stopped. This value is ignored if `enableDiskOffload` is `false`.
- **Use listener context (Property name: `useListenerContext`, Default: `false`).** Used for cache invalidation listener. If the user wants to use listener J2EE context for callback, set this to `true`. If the user wants to use the caller thread context for callback, set this to `false`.

There are three methods for configuring and using cache instances.

### Method 1 - Administrative Console

You can create additional cache instances using the Administrative Console. Click **Resources > Cache Instances**. After you create a new cache instance, you can optionally add additional properties, such as `flushToDiskOnStop` and `useListenerContext`, using the custom properties. See [Cache instance settings](#) for more information.

If you defined two cache instances in the Administrative Console with JNDI names of **services/cache/instance\_one** and **services/cache/instance\_two**, you can use the following code to look up the cache instances:

```
InitialContext ic = new InitialContext();
    DistributedMap dm1 = (DistributedMap)ic.lookup("services/cache/instance_one");
    DistributedMap dm2 = (DistributedMap)ic.lookup("services/cache/instance_two");
```

**Note:** For more information about the DistributedMap interface, see the Javadoc ([../javadoc/ee/com/ibm/websphere/cache/DistributedMap.html](http://../javadoc/ee/com/ibm/websphere/cache/DistributedMap.html)) for the `com.ibm.websphere.cache` package.

### Method 2 - Properties file

**Note:** The `distributedmap.properties` file will be deprecated in a future release. Using the Administrative Console is the recommended way to configure cache instances.

You can create additional cache instances using the `distributedmap.properties` file with the following format. :

```
cache.instance.0=/services/cache/instance_one
cache.instance.0.cacheSize=1000
cache.instance.0.enableDiskOffload=true
cache.instance.0.diskOffloadLocation=${WAS_INSTALL_ROOT}/diskOffload
cache.instance.0.flushToDiskOnStop=true
cache.instance.0.useListenerContext=true
cache.instance.1=/services/cache/instance_two
cache.instance.1.cacheSize=1500
cache.instance.1.enableDiskOffload=false
cache.instance.1.flushToDiskOnStop=false
cache.instance.1.useListenerContext=false
```

The preceding example creates two DistributedMap instances named `instance_one` and `instance_two`. `instance_one` has a cache entry size of 1,000 and `instance_two` has a cache entry size of 1,500. Disk offload is enabled in `instance_one` and disabled in `instance_two`. Use listener context is enabled in `instance_one` and disabled in `instance_two`. Flush to disk on stop is enabled in `instance_one` and disabled in `instance_two`.

You must place the `distributedmap.properties` file in either your application server or application class path. For example, you can use your application WAR file, `WEB-INF\classes` directory, or `was_root\classes` directory. The first entry in the properties file (`cache.instance.0`) specifies the JNDI name for the cache instance in the global namespace. You can use the following code to look up the cache instance:

```
InitialContext ic = new InitialContext();
    DistributedMap dm1 = (DistributedMap)ic.lookup("services/cache/instance_one");
    DistributedMap dm2 = (DistributedMap)ic.lookup("services/cache/instance_two");
```

**Note:** For more information about the DistributedMap interface, see the Javadoc ([../javadoc/ee/com/ibm/websphere/cache/DistributedMap.html](http://../javadoc/ee/com/ibm/websphere/cache/DistributedMap.html)) for the `com.ibm.websphere.cache` package.

### Method 3 - Resource references

**Note:** Method three is an extension to method one or method two, listed above. First use either method one or method two.

The following describes resource reference files. Alternately, you can use the Application Assembly Tool to create the resource reference.

Define a resource-ref in your module deployment descriptor (web.xml and ibm-web-bnd.xmi files) and look up the cache using the java:comp namespace.

**Resource-ref example:**

**File: web.xml**

```
<resource-ref id="ResourceRef_1">
  <res-ref-name>dmap/LayoutCache</res-ref-name>
  <res-type>com.ibm.websphere.cache.DistributedMap</res-type>
  <res-auth>Container</res-auth>
  <res-sharing-scope>Shareable</res-sharing-scope>
</resource-ref>
<resource-ref id="ResourceRef_2">
  <res-ref-name>dmap/UserCache</res-ref-name>
  <res-type>com.ibm.websphere.cache.DistributedMap</res-type>
  <res-auth>Container</res-auth>
  <res-sharing-scope>Shareable</res-sharing-scope>
</resource-ref>
```

**File: ibm-web-bnd.xmi**

```
<?xml version="1.0" encoding="UTF-8"?>
<webappbnd:WebAppBinding xmi:version="2.0" xmlns:xmi="http://www.omg.org/XMI"
  xmlns:webappbnd="webappbnd.xmi"
  xmlns:webapplication="webapplication.xmi" xmlns:commonbnd="commonbnd.xmi" xmlns:common="common.xmi"
  xmi:id="WebApp_ID_Bnd" virtualHostName="default_host">
  <webapp href="WEB-INF/web.xml#WebApp_ID"/>
  <resRefBindings xmi:id="ResourceRefBinding_1" jndiName="services/cache/instance_one">
    <bindingResourceRef href="WEB-INF/web.xml#ResourceRef_1"/>
  </resRefBindings>
  <resRefBindings xmi:id="ResourceRefBinding_2" jndiName="services/cache/instance_two">
    <bindingResourceRef href="WEB-INF/web.xml#ResourceRef_2"/>
  </resRefBindings>
</webappbnd:WebAppBinding>
```

The following example shows how to look up the resource-ref:

```
InitialContext ic = new InitialContext();
DistributedMap dm1a = (DistributedMap)ic.lookup("java:comp/env/dmap/LayoutCache");
DistributedMap dm2a = (DistributedMap)ic.lookup("java:comp/env/dmap/UserCache");
```

The previous resource-ref example maps java:comp/env/dmap/LayoutCache to /services/cache/instance\_one and java:comp/env/dmap/UserCache to /services/cache/instance\_two. In the examples, DistributedMap dm1 and dm1a are the same map. DistributedMap dm2 and dm2a are the same map.

**Note:** The DistributedMap interface does not currently have authorization or access control associated with any of the cache entries.

### What to do next

To learn how to share cached objects in a clustered environment, see [Sharing cached objects in a clustered environment](#).

## Sharing cached objects in a clustered environment

In a clustered environment, the content you place in cache might be shared with other servers in the cluster. The content might also be off-loaded to disk. If you intend to have the cached objects shared or off-loaded to disk, you must make these particular objects serializable. If the objects you place in cache are non-serializable, you must specify that the sharing policy for these objects is "not shared". The DistributedMap interface Javadoc

(../javadoc/ee/com/ibm/websphere/cache/DistributedMap.html) contains information about how to specify the sharing policy for a cached object. Specifying a sharing policy other than "not shared" for non-serializable objects can result in poor system performance.

## Cache instance settings

A cache instance is a location, in addition to the default shared dynamic cache, where any J2EE application can store, distribute, and share data. This gives applications greater flexibility and better tuning of the cache resources. The programming interface to access this cache instance is called DistributedMap. See the DistributedMap API documentation in the WebSphere Application Server Javadoc (../javadoc/ee/com/ibm/websphere/cache/DistributedMap.html) for more information.

To view this administrative console page, click **Resources > Cache Instances**.

### Name

Specifies the required display name for the resource.

### JNDI name

Specifies the Java Naming and Directory Interface (JNDI) name for the resource. Use this name when looking up a reference to this cache instance. The returned object is of type DistributedMap.

### Description

Specifies a description for the resource. This field is optional.

### Category

Specifies a category string to classify or group the resource. This field is optional.

### Cache Size

Specifies a positive integer for the maximum number of entries the cache holds.

The range is 100 to 200,000.

### Default Priority

Specifies the default priority for servlets that can be cached. This value will determine how long an entry will stay in a full cache.

The recommended value is one.

### Enable Disk Offload

Checking this box will enable disk offloading.

### Disk Offload Location

Specifies the directory used for disk offload.

### Additional custom properties (available only when a cache instance is created)

Custom properties are available for flushToDiskOnStop and useListenerContext.

**Note:** You can define custom properties once you create a cache instance.

- Name - Specifies the name of the custom property. In this case, it will be flushToDiskOnStop or useListenerContext. This field is required.
- Value - True or false. Specifies the value of the custom property. This field is required.

- **Description** - Specifies text to describe any bounds or values for this property. This field is optional.

## Cache instance collection

Use this page to configure and manage cache instances.

To view this administrative console page, click **Resources > Cache Instances > New**.

### Scope

Specify **CELL SCOPE** to view and configure cache instances that are available to all servers within the cell. Specify **NODE SCOPE** to view and configure cache instances that are available to all servers with the particular node. Specify **SERVER SCOPE** to view and configure cache instances that are available only on the specific server.

### Name

Specifies the required display name for the resource.

### JNDI name

Specifies the Java Naming and Directory Interface (JNDI) name for the resource. Use this name when looking up a reference to this cache instance. The returned object is of type `DistributedMap`.

### Description

Specifies a description for the resource. This field is optional.

### Category

Specifies a category string to classify or group the resource. This field is optional.

### Cache Size

Specifies a positive integer for the maximum number of entries the cache holds.

The range is 100 to 200,000.

### Enable Disk Offload

Checking this box will enable disk offloading.

By default, dynamic cache only maintains the number of entries configured in memory. If new entries are created while the cache is full, the priorities configured for each cache entry, along with a least recently used (LRU) algorithm, are used to remove entries from the cache. As an alternative, you can configure disk offload to copy entries onto the file system to be accessed at a later time. The location to copy the entries is configurable.

## Cache instance service settings

The Cache Instance Service enables server-level configuration for Cache Instance resources.

To view this administrative console page, click **Servers > Application Servers > server1 > Cache Instance Service**.

### Startup

Specifies whether the server will attempt to start the specified service when the server starts. Checking the box turns this service on or off.

## Invalidation listeners

Invalidation listener mechanism uses Java events for alerting applications when contents are removed from the cache.

Applications implement the `InvalidationListener` interface (defined in the `com.ibm.websphere.cache` package) and register it to the cache using the `DistributedMap` interface. Listeners receive `InvalidationEvents` (defined in the `com.ibm.websphere.cache` package) when entries from the cache are evicted, due to an explicit user invalidation, timeout, least recently used (LRU) eviction, cache clear, or disk timeout. Applications can immediately recalculate the invalidated data and prime the cache before the next user request.

Enable listener support in `DistributedMap` before registering listeners. `DistributedMap` can also be configured to use the invalidation listener J2EE context from registration time during callbacks. Setting the value of the custom property `useListenerContext` to `true` will enable the invalidation listener J2EE context for callbacks. See Cache instance settings for more information.

The following example shows how to set up an invalidation listener:

```
dmap.enableListener(true); // Enable cache invalidation listener.
InvalidationListener listener = new MyListenerImpl();
    //Create invalidation listener object.
dmap.addInvalidationListener(listener); //Add invalidation listener.
    :
    :
    :
dmap.removeInvalidationListener(listener);
    //Remove the invalidation listener. This increases performance.
dmap.enableListener(false);
    // Disable cache invalidation listener. This increases performance.
```

For more information about invalidation listeners, see the Javadoc ([../javadoc/ee/com/ibm/websphere/cache/package-summary.html](http://www.ibm.com/websphere/cache/package-summary.html)) for the `com.ibm.websphere.cache` package.

---

## Example: Caching Web services

The following is an example of building a set of cache policies for a simple Web services application. The application in this example stores stock quotes, and has operations to read, update the price of, and buy a given stock symbol.

Following are two SOAP message examples that the application can receive, with accompanying HTTP Request headers.

The first message sample contains a SOAP message for a `GetQuote` operation, requesting a quote for IBM. This is a read-only operation that gets its data from the back-end, and is very cacheable. In this example the SOAP message is cached and a timeout is placed on its entries to guarantee the quotes it returns are not too out of date.

### Message example 1

```
POST /soap/servlet/soaprouter
HTTP/1.1
Host: www.myhost.com
Content-Type: text/xml; charset="utf-8"
SOAPAction: urn:stockquote-lookup
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
```

```

<SOAP-ENV:Body>
<m:getQuote xmlns:m="urn:stockquote:>
<symbol>IBM</symbol>
</m:getQuote>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

The SOAPAction HTTP header in the request is defined in the SOAP specification and is used by HTTP proxy servers to dispatch requests to particular HTTP servers. WebSphere Application Server dynamic cache can use this header in its cache policies to build IDs without having to parse the SOAP message.

Message example 2 illustrates a SOAP message for a BuyQuote operation. While message 1 is cacheable, this message is not, because it updates the back-end database.


### Message example 2

```

POST /soap/servlet/soaprouter
HTTP/1.1
Host: www.myhost.com
Content-Type: text/xml; charset="utf-8"
SOAPAction: urn:stockquote-update
<SOAP-ENV:Envelope xmlns:SOAP-ENV="http://schemas.xmlsoap.org/soap/envelope/"
SOAP-ENV:encodingStyle="http://schemas.xmlsoap.org/soap/encoding/">
<SOAP-ENV:Body>
<m:buyStock xmlns:m="urn:stockquote:>
<symbol>IBM</symbol>
</m:buyStock>
</SOAP-ENV:Body>
</SOAP-ENV:Envelope>

```

The graphic illustrates how to invoke methods with the SOAP messages. In Web services terms, especially Web Service Definition Language (WSDL), a service is a collection of operations such as getQuote and buyStock. A body element namespace (urn:stockquote in our example) defines a service, and the name of the first body element indicates the operation.



The following is an example of WSDL for the getQuote operation:

```

<?xml version="1.0"?>
<definitions name="StockQuoteService-interface"
targetNamespace="http://www.getquote.com/StockQuoteService-interface"
xmlns:tns="http://www.getquote.com/StockQuoteService-interface"
xmlns:xsd="http://www.w3.org/2001/XMLSchema"
xmlns:soap="http://schemas.xmlsoap.org/wsdl/soap/"
xmlns="http://schemas.xmlsoap.org/wsdl/"
<message name="SymbolRequest">

```

```

<part name="return" type="xsd:string"/>
</message>
<portType name="StockQuoteService">
<operation name="getQuote">
<input message="tns:SymbolRequest"/>
<output message="tns:QuoteResponse"/>
</operation>
</portType>
<binding name="StockQuoteServiceBinding"
type="tns:StockQuoteService">
<soap:binding style="rpc" transport="http://schemas.xmlsoap.org/soap/http"/>
<operation name="getQuote">
<soap:operation soapAction="urn:stockquote-lookup"/>
<input>
<soap:body use="encoded" namespace="urn:stockquote"
encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
</input>
<output>
<soap:body use="encoded" namespace="urn:stockquotes"
encodingStyle="http://schemas.xmlsoap.org/soap/encoding"/>
</output>
</operation>>
</binding>
</definition>

```

To build a set of cache policies for a Web services application configure WebSphere Application Server dynamic cache to recognize cacheable service operation of the operation.

WebSphere Application Server inspects the HTTP request to determine whether or not an incoming message can be cached based on the cache policies defined for an application. In this example, buyStock and stock-update are not cached, but stockquote-lookup is cached. In the cachespec.xml file for this Web application, the cache policies need defining for these services so that the dynamic cache can handle both SOAPAction and service operation.

WebSphere Application Server uses the operation and the message body in Web services cache IDs, each of which has a component associated with them. Therefore, each Web services <cache-id> rule contains only two components. The first is for the operation. Because you can perform the stockquote-lookup operation by either using a SOAPAction header or a service operation in the body, you must define two different <cache-id> elements, one for each method. The second component is of type "body", and defines how WebSphere Application Server should incorporate the message body into the cache ID. You can use a hash of the body, although it is legal to use the literal incoming message in the ID.

The incoming HTTP request is analyzed by WebSphere Application Server to determine which of the <cache-id> rules match. Then, the rules are applied to form cache or invalidation IDs.

The following is sample code of a cachespec.xml file defining SOAPAction and servicesOperation rules:

```

<cache>
<cache-entry>
<class>webservice</class>
<name>/soap/servlet/soaprouter</name>
<sharing-policy>not-shared</sharing-policy>
<cache-id>
<component id="" type=SOAPAction>
<value>urn:stockquote-lookup</value>
</component>

```



```

<component id="Hash" type="SOAPEnvelope"/>
<timeout>3600</timeout>
<priority>1</priority>
</cache-id>
<cache-id>
<component id="" type="serviceOperation">
<value>urn:stockquote:getQuote</value>
</component>
<component id="Hash" type="SOAPEnvelope"/>
<timeout>3600</timeout>
<priority>1</priority>
</cache-id>
</cache-entry>
</cache>

```

---

## Example: Configuring the dynamic cache

This example puts all the steps together for configuring the dynamic cache with the `cachespec.xml` file, showing the use of the cache ID generation rules, dependency IDs, and invalidation rules.

Suppose we have a servlet which is used to manage a simple news site. This servlet uses the query parameter "action" to determine whether the request is being used to "view" news or "update" news (used by the administrator). Further, another query parameter "category" is used to select the news category. Further, suppose that this site supports an optional customized layout, which is stored in the user's session using the attribute name "layout". Here are example URL requests to this servlet:

`http://yourhost/yourwebapp/newscontroller?action=view&category=sports` (Returns a news page for the sports category )

`http://yourhost/yourwebapp/newscontroller?action=view&category=money` (Returns a news page for the money category)

`http://yourhost/yourwebapp/newscontroller?action=update&category=fashion`  
(Allows the administrator to update news in the fashion category)

Here are the steps for configuring dynamic cache with `cachespec.xml`, using the information provided to you:

1. Define the cache-entry elements necessary to identify the servlet. In this case, the servlet's URI is "newscontroller" so this will be our cache-entry's name element. Also, since we are caching a servlet/JavaServer Page (JSP), the cache-entry class is "servlet".

```

<cache-entry>
<name> /newscontroller </name>
<class>servlet </class>
</cache-entry>

```

2. Define cache ID generation rules. For this servlet, we only want to cache when `action=view`, so one component of the cache ID will be the parameter "action" when the value equals "view". The news category is also an essential part of the cache ID. Finally, the optional session attribute for the user's layout is included in the cache ID. The cache-entry now looks like this:

```

<cache-entry>
<name> /newscontroller </name>
<class>servlet </class>
<cache-id>
<component id="action" type="parameter">
<value>view</value>

```

```

        <required>true</required>
      </component>
      <component id="category" type="parameter">
        <required>true</required>
      </component>
      <component id="layout" type="session">
        <required>>false</required>
      </component>
    </cache-id>
  </cache-entry>

```

3. Define dependency ID rules. For this servlet, a dependency ID will be added for the category. Later, when the category is invalidated due to an update event, all views of that news category will be invalidated. After adding our dependency-id, the cache-entry now looks like this:

```

<cache-entry>
  <name>newscontroller </name>
  <class>servlet </class>
  <cache-id>
    <component id="action" type="parameter">
      <value>view</value>
      <required>true</required>
    </component>
    <component id="category" type="parameter">
      <required>true</required>
    </component>
    <component id="layout" type="session">
      <required>>false</required>
    </component>
  </cache-id>
  <dependency-id>category
    <component id="category" type="parameter">
      <required>true</required>
    </component>
  </dependency-id>
</cache-entry>

```

4. Define invalidation rules. Since we defined a category dependency ID, we will now define an invalidation rule to invalidate the category when action=update. To incorporate the conditional logic, we will add "ignore-value" components into the invalidation rule. These components will not add to the output of the invalidation ID, but will only determine whether or not the invalidation ID is created and executed. The final cache-entry now looks like this:

```

<cache-entry>
  <name>newscontroller </name>
  <class>servlet </class>
  <cache-id>
    <component id="action" type="parameter">
      <value>view</value>
      <required>true</required>
    </component>
    <component id="category" type="parameter">
      <required>true</required>
    </component>
    <component id="layout" type="session">
      <required>>false</required>
    </component>
  </cache-id>
  <dependency-id>category
    <component id="category" type="parameter">
      <required>true</required>
    </component>
  </dependency-id>
  <invalidation>category<
    <component id="action" type="parameter" ignore-value="true">
      <value>update</value>

```

```
        <required>true</required>
    </component>
    <component id="category" type="parameter">
        <required>true</required>
    </invalidation>
</cache-entry>
```

---

## Cache monitor

Cache monitor is an installable Web application that provides a real time and run time view of the current state of the cache, and helps verify that the cache is behaving as expected. The only way to manipulate the data in the cache is by using the cache monitor. It provides a GUI interface to manually change data.

Cache monitor does not provide a view for cache instance resources.

Cache monitor provides a way to:

- **Verify the configuration of dynamic cache**

The WebSphere Application Server Administration Console provides ways to enable the dynamic cache service and configure properties, such as maximum size of the cache and disk offload location, as well as advanced features such as controlling external caches. Cache monitor offers a way for dynamic cache users to verify the configuration of the dynamic cache by providing a convenient view of the configured features and properties in the cache monitor.

- **Verify the cache policies**

To cache an object, WebSphere Application Server must know how to generate unique IDs for different invocations of that object. This is performed by providing rules for each cacheable object in the `cachespec.xml` file, found inside the Web module `WEB-INF` or enterprise bean `META-INF` directory. Each cacheable object can have multiple cache ID rules that execute in sequence until either a rule returns a cache ID or no more rules remain to execute. If none of the cache ID generation rules produce a valid cache ID, then the object is not cached. Since there can be multiple `cachespec.xml` files with multiple cache ID rules, cache monitor provides a convenient way to verify the policies of each object. It offers a view of all the cache policies currently loaded in dynamic cache. This view is also convenient to verify that the `cachespec.xml` file was read by the dynamic cache without errors.

- **Monitor cache statistics**

Cache monitor provides a view of the essential cache data, such as number of cache hits, cache misses, and number of entries in cache. This helps to tune the cache configuration optimally to get the best performance improvement out of dynamic cache. For example, if the number of used entries is often high, and entries are being removed and recreated, one might consider increasing the maximum size of the cache or enabling disk offload.

- **Monitor the data flowing through the cache**

Once a cacheable object is invoked, dynamic cache creates a cache entry for it that contains the output of the execution and metadata, such as time to live, sharing policy, etc. Entries are distinguished by a unique ID string that is based on the rules specified in the `cachespec.xml` file for this object's name. Objects with the same name may generate multiple cache IDs for different invocations, based on request parameters and attributes for each invocation. Cache monitor provides a view of all the cache entries currently in cache, based on the unique ID. It also provides a view of the group of cache entries that share a common name (also known as template). Cache entries can also be grouped together by a dependency ID, which is used to invalidate the entire group of entries

dependent on a common entity. Therefore, cache monitor also provides a view of the group of cache entries that share a common dependency ID.

For each entry, cache monitor also displays metadata, such as time to live, priority and sharing-policy, and provides a view of the output that has been cached. This helps the customer to verify which pages have been cached, that the pages have been cached with the right attributes such as time to live, priority, etc., and that the pages have the right content.

- **Monitor the data in the edge cache**

Dynamic cache provides support to recognize the presence of an Edge Side Include (ESI) processor and to generate ESI include tags and appropriate cache policies for edge cacheable fragments. The ESI processor has the ability to cache whole pages, as well as fragments, providing a higher cache hit ratio. There can be multiple ESI processors running on multiple hosts configured for caching.

Cache monitor provides a list of all ESI processes and their hosts that are enabled for caching. It also provides a way to select a host or a processor, and view its edge cache statistics as well as current cache entries.

- **View the data offloaded to the disk**

By default, when the number of cache entries reaches the configured limit for a given server, eviction of cache entries occurs, allowing new entries to enter the cache service. The dynamic cache includes the disk offload feature that copies the evicted cache entries to disk for future access. Cache monitor offers a view of the content offloaded to disk that corresponds to the view of contents cached in memory.

- **Manage the data in the cache**

Besides displaying cache content, cache monitor also provides some basic operations on the data in the cache:

- Removing an entry from the cache
- Removing all entries for a certain dependency ID
- Removing all entries for a certain name (template)
- Moving an entry to the front of the least recently used queue to avoid eviction
- Moving an entry from the disk to the cache
- Clearing the entire contents of the cache
- Clearing the contents of the disk cache

These functions are useful for dynamic cache customers, as they provide a way to manually change the state of the cache without having to restart the server.

## Edge cache statistics

Cache monitor provides a view of the edge cache statistics.

The following statistics are available:

- **ESI Processors.** Number of processes configured as edge caches.
- **Number of Edge Cached Entries.** Number of entries currently cached on all edge servers and processes.
- **Cache Hits.** Number of requests that match entries on edge servers.
- **Cache Misses By URL.** A cache policy does not exist on the edge server, for the requested template.

**Note:**

- The initial ESI request for a template that has a cache policy on WebSphere Application Server will result in a miss.
- Every request for a template that does not have a cache policy on WebSphere Application Server will result in a miss by URL on the edge server.
- **Cache Misses By Cache ID.** The policy for the requested template exists on the edge server, and a cache ID is created, based on the ID rules and the request attributes, but the cache entry for this ID does not exist.  
**Note:** If the policy exists on the edge server for the requested template, but a cache ID match is not found, based on the ID rules and the request attributes, it is not treated as a cache miss.
- **Cache Timeouts.** Number of entries removed from the edge cache, based on the timeout value.
- **Evictions.** Number of entries removed from the edge cache, due to invalidations received from WebSphere Application Server.



---

## Chapter 26. Managing user profiles

**Note:** User Profile Manager API is deprecated in the current release, and there is no replacement available.

IBM WebSphere Application Server provides a service for processing user profiles, called the *User Profile Manager*. The service is provided in the form of an EJB entity bean that servlets can call whenever they are required to access a user profile.

The key activities for implementing user profiles are summarized.

### Steps for this task

1. Customize the user profile support as necessary.

Options include:

- Using the data representation class with the name-value pairs it currently supports (no action required)
- Extending the data representation class to support additional, arbitrary name-value pairs
- Adding columns to the base user profile representation
- Extending the User Profile enterprise bean to import existing databases

Evaluate whether the user profile representation provided by IBM represents the kind of data you want to keep about your users. You might find it desirable to customize the IBM user profile support.

2. Create or modify servlets to use the User Profile Manager and related user profile support classes to maintain user profiles on behalf of Web applications.
3. Assemble your application.
4. Deploy your application.
5. Ensure the administrator appropriately configures User Profile Managers using `userprofile.xml` file.

If the programmer and administrator are not the same person, the programmer might need to provide settings information to the administrator, based on how the programmer implemented user profiles.

---

## User profile

Some applications collect data about the users with which they interact. The data is stored in a database. The next time the user interacts with the application, the application recalls the data.

Because the application already knows something about the user, it can provide the user with a more personalized experience.

User profiles provide a means by which a company can maintain and manage database tables containing fields for demographic data and use those tables to interact with a database of individual customers or other users on the company system.

For example, when a repeat user logs onto a Web site that supports user profiles, the Web site can display headlines and advertising tailored to the shopping preferences of that user. The site can address the user by logon name.

An application implementing user profiles requires database access for storing the user profile data it gathers.

---

## UserProfileManager class

Servlets and other application building blocks requiring user profile support should make calls to the class:

```
com.ibm.websphere.userprofile.UserProfileManager
```

The class supports the following functions:

- Creating and deleting user profiles
- Getting and updating (cached and immediate) to and from the database
- Getting user profiles for read-only tasks
- Performing queries on database columns

---

## User profile development options

The application developer has a few options for customizing the user profile support provided by IBM WebSphere Application Server. The Related information provides instructions and additional details about each option.

### Extending the data represented in user profiles

Web applications can maintain several pieces of data about users. You can extend the data representation to allow the collection of arbitrary name-value pairs.

Use the following interface with the `com.ibm.websphere.userprofile.UserProfileExtender` class to extend a user profile hash table:

```
com.ibm.websphere.userprofile.UserProfileProperties
```

This action enables you to place arbitrary name-value pairs in the user profile. Extending the hash table is similar to using the `java.util.Dictionary` class in the base JDK 1.x, or any of the classes that extend it.

### Adding columns to the base user profile implementation

Application developers can customize user profiles by adding columns to the base user profile implementation. Adding new columns is accomplished by implementing the interface:

```
com.ibm.websphere.userprofile.UserProfileExtender
```

and extending the base class:

```
com.ibm.servlet.personalization.userprofile.UserProfile
```

The application developer can add columns to but not delete columns from the base implementation.

Adding columns is a two-step process, as follows:

1. Extend the `UserProfile` class.
2. Modify your existing servlets to use the new columns.



Several examples are available to demonstrate how to extend the base user profile implementation and utilize the extension with a servlet.

| Example                        | Description  |
|--------------------------------|--|
| UPServletExample.java          | Demonstrates how a servlet opens a user profile and prints the fields contained within.  |
| UserProfileExtendedSample.java | Shows how to extend the UserProfile class to add a column to the user profile for a cellular phone number.<br><br>The WebSphere Application Server administrator configures the User Profile Manager to point to the extended class. |
| UPServletExampleExtended.java  | Shows how to modify the UPServletExample servlet to include the cellular phone number in the output.   |
| UserProfileExtended.java       | Shows how to extend a hash table to place arbitrary name-value pairs into the user profile.  |
| UPServletExtended.java         | Shows how to extend the servlet. When any of the newly added columns are removed or replaced, look for the table named "USERPROFILE" in the database to which the user profile is configured and drop that table.                    |

The examples are encoded in HTML for viewing in a browser. The documentation directory also contains nonHTML versions (.java files) that are ready for use.

## Extending the User Profile enterprise bean and importing legacy databases

Application developers can extend the User Profile enterprise bean itself and import legacy databases into the user profile. The main advantage in extending the User Profile enterprise bean is to gain the ability to import existing databases into the user profile. You can also extend this enterprise bean to add columns to the base user profile implementation.

### UPServletExample.java

```
import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;

//Creates a Userprofile using the new API

public class UPServlet_ReadWrite extends HttpServlet {

    public void doGet (HttpServletRequest req, HttpServletResponse res)
        throws ServletException, IOException {
```

```

PrintWriter out;
res.setContentType("text/html");
out = res.getWriter();

UserProfileManager manager = UserProfileManager.getUserProfileManager();
UserProfile userprofile;

try {

    //Try creating the UserProfile
    userprofile = manager.addUserProfile("bpink");

} catch(UserProfileCreateException e1) {

    try {

        //Try finding the existing in readWrite mode.
        //Second argument indicates whether we want to get userprofile
        //in read only mode or read write mode.

        userprofile = manager.getUserProfile("bpink",true);

    } catch(UserProfileFinderException e) {
        e.printStackTrace();
        return;
    }

}

//Set the properties

userprofile.setAddress1("myaddress1");
userprofile.setAddress2("myaddress2");
userprofile.setFirstName("Pinkowski");
userprofile.setSurName("Ben");
userprofile.setDayPhone("555-6677");
userprofile.setNightPhone("556-6765");
userprofile.setCity("MYCITY");
userprofile.setNation("myCountry");
userprofile.setEmployer("MyEmployer");
userprofile.setFax("7823470");

userprofile.setLanguage("mylanguage");
userprofile.setEmail("MyEmail@email");
userprofile.setStateOrProvince("myState");
userprofile.setPostalCode("xxxxx");

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//Checking whether it updated the info

try {

    //Getting the existing userprofile in ReadOnly mode.

    userprofile = manager.getUserProfile("bpink",false);

} catch(UserProfileFinderException e1) {

    out.println("Error finding ");
}

```

```

        e1.printStackTrace();
        return;
    }

    //Displaying the properties of userprofile

    out.println(userprofile.getAddress1()+"<br>");
    out.println(userprofile.getAddress2()+"<br>");;
    out.println(userprofile.getFirstName()+"<br>");;
    out.println(userprofile.getSurName()+"<br>");;
    out.println(userprofile.getDayPhone()+"<br>");;
    out.println(userprofile.getNightPhone()+"<br>");;
    out.println(userprofile.getCity()+"<br>");;
    out.println(userprofile.getNation()+"<br>");;
    out.println(userprofile.getEmployer()+"<br>");;
    out.println(userprofile.getFax()+"<br>");;
    out.println(userprofile.getLanguage()+"<br>");;
    out.println(userprofile.getEmail()+"<br>");;
    out.println(userprofile.getStateOrProvince()+"<br>");;
    out.println(userprofile.getPostalCode()+"<br>");;

    //Freeing resources held by userprofile
    manager.releaseResources(userprofile);
}
}

```

## UserProfileExtendedSample.java

```

/* -----
** Copyright 1997-99 IBM Corporation. All rights reserved.
**
** -----
*/
package com.ibm.servlet.personalization.userprofile;

import com.ibm.servlet.personalization.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileExtender;

//Extensions of UserProfile to add new Columns should implement UserProfileExtender
public class UserProfileExtendedSample
    extends com.ibm.servlet.personalization.userprofile.UserProfile
    implements UserProfileExtender {

    //New column that is being added by this
    //derived class.
    public String cellPhone;

    //Manager Class will call this method to append new Column types.
    //If UserProfile class is extended to append new columns
    //TOTAL COLUMNS: Base Class columns + columns returned by this class

    public String[] getNewColumns() {
        //If variable name is "cellPhone," you need to
        //return "cellPhone" in array format. JDBC equivalent will be
        //generated automatically. You can add multiple columns.
        //For multiple columns: String newCol={"fieldName1","fieldName2",...};
        String[] newCol={"cellPhone"};
        return newCol;
    }

    public String getCellPhone() {
        // Need to call this method to
        // get the things from persistence store.
        return(String)getByType("cellPhone");
    }
}

```

```

    }

    public void setCellPhone(String value) {
        cellPhone = value;
        //Call this method to store the
        //things in persistence store
        setByType("cellPhone", value);
    }
}

```

## UPServletExampleExtended.java

```

import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;
import com.ibm.servlet.personalization.userprofile.UserProfileExtendedSample;

public class UPServletExtendedSample extends HttpServlet {

    public void doGet (HttpServletRequest req, HttpServletResponse res)
    throws ServletException, IOException
    {
        UserProfileManager manager = UserProfileManager.getUserProfileManager();
        UserProfile userprofile;

        PrintWriter out;

        res.setContentType("text/html");
        out = res.getWriter();

        try {

            //try Creating the UserProfile

            userprofile = manager.addUserProfile("bpink");

        } catch( UserProfileCreateException e1) {

            try { //try finding the existing in readWrite mode

                userprofile = manager.getUserProfile("bpink",true);
            } catch(UserProfileFinderException e) {
                e.printStackTrace();
                return;
            }

        }

        userprofile.setAddress1("myaddress1");
        userprofile.setAddress2("myaddress2");
        userprofile.setFirstName("Pinkowski");
        userprofile.setSurName("Ben");
        userprofile.setDayPhone("555-6677");
        userprofile.setNightPhone("556-6765");
    }
}

```

```

userprofile.setCity("MYCITY");
userprofile.setNation("myCountry");
userprofile.setEmployer("MyEmployer");
userprofile.setFax("7823470");
userprofile.setLanguage("mylanguage");
userprofile.setEmail("MyEmail@email");
userprofile.setStateOrProvince("myState");
userprofile.setPostalCode("xxxxx");

//calling setCellPhone
((com.ibm.servlet.personalization.userprofile.UserProfileExtendedSample)
  userprofile).setCellPhone("346-4588");

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//Checking whether it updated the info
try {

    //Getting the existing userprofile
    userprofile = manager.getUserProfile("bpink",false);

} catch( UserProfileFinderException e1) {

    out.println("Error finding ");
    e1.printStackTrace();
    return;
}

out.println(userprofile.getAddress1()+"<br>");
out.println(userprofile.getAddress2()+"<br>");
out.println(userprofile.getFirstName()+"<br>");
out.println(userprofile.getSurName()+"<br>");
out.println(userprofile.getDayPhone()+"<br>");
out.println(userprofile.getNightPhone()+"<br>");
out.println(userprofile.getCity()+"<br>");
out.println(userprofile.getNation()+"<br>");
out.println(userprofile.getEmployer()+"<br>");
out.println(userprofile.getFax()+"<br>");
out.println(userprofile.getLanguage()+"<br>");
out.println(userprofile.getEmail()+"<br>");
out.println(userprofile.getStateOrProvince()+"<br>");
out.println(userprofile.getPostalCode()+"<br>");

//Calling getCellPhone
out.println(((UserProfileExtendedSample)userprofile).getCellPhone()+"<br>");

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//For getting values by cellPhone
out.println("<br><br>Retreiving by Cell Phone <br>");
Enumeration enum = manager.findUserProfiles("cellPhone","346-4588");
while(enum.hasMoreElements()) {

    com.ibm.websphere.userprofile.UserProfile up =
        (com.ibm.websphere.userprofile.UserProfile)enum.nextElement();
    out.println("first name :"+up.getFirstName()+"<br>");
}

```

```

        //Freeing resources held by userprofile
        manager.releaseResources(up);
    }
}

```

## UserProfileExtended.java

```

package com.ibm.servlet.personalization.userprofile;
/* -----
** Copyright 1997-99 IBM Corporation. All rights reserved.
** -----
*/
import java.util.*;

import com.ibm.servlet.personalization.userprofile.UserProfile;

import com.ibm.websphere.userprofile.UserProfileExtender;
import com.ibm.websphere.userprofile.UserProfileProperties;

public class UserProfileExtended extends UserProfile implements UserProfileExtender,
    UserProfileProperties {
    //New column that is being added by this
    //derived class.
    public Hashtable properties;

    static String propCol ="properties";

    //Manager Class will call this method to append new Column types
    //to SQL Strings. If UserProfile class is extended to append new columns
    //it should implement UserProfileExtender.
    //COLUMNS: Base Class columns + columns returned by this class

    public String[] getNewColumns() {
        //if variable name is properties, you need to
        //return "properties" . JDBC equivalent will be
        //generated automatically.
        String[] newCol={propCol};
        return newCol;
    }

    public Object getValue(String key) {
        // Need to call this method to
        // get the things from persistent store
        properties = (Hashtable) getByType(propCol);

        if(properties != null)
            return properties.get(key);

        else return null;
    }

    public void putValue(String key, Object value) {

        properties =(Hashtable) getByType(propCol);

        if(properties == null)
            properties = new Hashtable();

        properties.put(key,value);

        //store in persistent store
        setByType(propCol, properties);
    }
}

```

```

        public void removeValue(String key) {
            properties = (Hashtable) getByType(propCol);

            if(properties == null)
                return;

            properties.remove(key);

            //store in persistent store
            setByType(propCol, properties);
        }
    }
}

```

## UPServletExtended.java

```

import java.io.*;
import java.util.*;

import javax.servlet.*;
import javax.servlet.http.*;

import com.ibm.websphere.userprofile.UserProfile;
import com.ibm.websphere.userprofile.UserProfileManager;
import com.ibm.websphere.userprofile.UserProfileCreateException;
import com.ibm.websphere.userprofile.UserProfileFinderException;
import com.ibm.websphere.userprofile.UserProfileRemoveException;
import com.ibm.websphere.userprofile.UserProfileProperties;

public class UPServletExtended extends HttpServlet {

    public void doGet (HttpServletRequest req, HttpServletResponse res)
        throws ServletException, IOException
    {
        UserProfileManager manager = UserProfileManager.getUserProfileManager();
        UserProfile userprofile;

        PrintWriter out;

        res.setContentType("text/html");
        out = res.getWriter();

        try {

            //try Creating the UserProfile

            userprofile = manager.addUserProfile("bpink");

        } catch( UserProfileCreateException e1) {

            try { //try finding the existing in readWrite mode

                userprofile = manager.getUserProfile("bpink",true);
            } catch(UserProfileFinderException e) {
                e.printStackTrace();
                return;
            }

        }

        userprofile.setAddress1("myaddress1");
        userprofile.setAddress2("myaddress2");
        userprofile.setFirstName("Pinkowski");
        userprofile.setSurName("Ben");
        userprofile.setDayPhone("555-6677");
    }
}

```

```

userprofile.setNightPhone("556-6765");
userprofile.setCity("MYCITY");
userprofile.setNation("myCountry");
userprofile.setEmployer("MyEmployer");
userprofile.setFax("7823470");
userprofile.setLanguage("mylanguage");
userprofile.setEmail("MyEmail@email");
userprofile.setStateOrProvince("myState");
userprofile.setPostalCode("xxxxx");

//calling putValue
((UserProfileProperties)userprofile).putValue("name","HHHHHHH");
((UserProfileProperties)userprofile).putValue("Date",new java.util.Date());

//Freeing resources held by userprofile
manager.releaseResources(userprofile);
userprofile=null;

//Checking whether it updated the info
try {

    //Getting the existing userprofile

    userprofile = manager.getUserProfile("bpink",false);

    out.println(userprofile.getAddress1()+"<br>");
    out.println(userprofile.getAddress2()+"<br>");;
    out.println(userprofile.getFirstName()+"<br>");;
    out.println(userprofile.getSurName()+"<br>");;
    out.println(userprofile.getDayPhone()+"<br>");;
    out.println(userprofile.getNightPhone()+"<br>");;
    out.println(userprofile.getCity()+"<br>");;
    out.println(userprofile.getNation()+"<br>");;
    out.println(userprofile.getEmployer()+"<br>");;
    out.println(userprofile.getFax()+"<br>");;
    out.println(userprofile.getLanguage()+"<br>");;
    out.println(userprofile.getEmail()+"<br>");;
    out.println(userprofile.getStateOrProvince()+"<br>");;
    out.println(userprofile.getPostalCode()+"<br>");;

    //Getting the values

    out.println(((UserProfileProperties)userprofile).getValue("name")+"<br>");
    out.println(((UserProfileProperties)userprofile).getValue("Date")+"<br>");
    out.println("Removing Values ");
    ((UserProfileProperties)userprofile).removeValue("name");
    ((UserProfileProperties)userprofile).removeValue("Date");
    out.println(((UserProfileProperties)userprofile).getValue("name")+"<br>");
    out.println(((UserProfileProperties)userprofile).getValue("Date")+"<br>");

    //Freeing resources held by userprofile
    manager.releaseResources(userprofile);

} catch( UserProfileFinderException e1) {

    out.println("Error finding ");
    e1.printStackTrace();
    return;
}
}
}

```



---

## userprofile.xml

To *installation\_root*/properties, add a file named userprofile.xml in the following format. Specify enterprise bean class names; data wrapper class name; and JNDI names for the read-only bean, read/write bean, and data source (from step 2). You must also add user ID and password information for the JNDI data source

The following example file contains class names as provided in WebSphere Application Server. If data wrapper and enterprise bean class names are extended programmatically, change them accordingly.

```
<?xml version="1.0"?>

  <userprofile>
    <userprofile-enabled>true</userprofile-enabled>
    <userprofile-wrapper-class>
      <classname>
        com.ibm.servlet.personalization.userprofile.UserProfile
      </classname>
    </userprofile-wrapper-class>
    <userprofile-manager-name>
      User Profile Manager
    </userprofile-manager-name>
    <userprofile-bean>
      <readonly-interface>
        com.ibm.servlet.personalization.userprofile.UP_ReadOnly
      </readonly-interface>
      <readwrite-interface>
        com.ibm.servlet.personalization.userprofile.UP_ReadWrite
      </readwrite-interface>
      <readonlyhome-interface>
        com.ibm.servlet.personalization.userprofile.UP_ReadOnlyHome
      </readonlyhome-interface>
      <readwritehome-interface>
        com.ibm.servlet.personalization.userprofile.UP_ReadWriteHome
      </readwritehome-interface>
      <readonly-JNDI-lookupName>UP_ReadOnlyHome</readonly-JNDI-lookupName>
      <readwrite-JNDI-lookupName>UP_ReadWriteHome</readwrite-JNDI-lookupName>
    </userprofile-bean>

    <userprofile-store>
      <database-userid></database-userid>
      <database-password></database-password>
      <database-datasource></database-datasource>
    </userprofile-store>
  </userprofile>
```



---

## Chapter 27. Assembling applications

Assemble application modules (known as EAR files) from new or existing J2EE 1.3 modules, including these archives: Web application archives (WAR), resource adapter archives (RAR), enterprise beans (EJB JAR), and application client archives (JAR). This packaging and configuration of code artifacts into application modules or standalone Web modules is necessary for deploying the applications onto the application server.

### Before you begin

Gather the code artifacts that you want to package into one or more assembled modules. Code artifacts include these items that you have created and unit tested in your favorite integrated development environment:

- Enterprise beans
- Servlets, JavaServer Pages (JSP) files and other Web components
- Resource adapter (*connector*) implementations
- Application clients
- Other supporting classes and files

### Steps for this task

1. Start the Application Assembly Tool (AAT).
2. **(Optional)** Migrate existing J2EE 1.2 modules to J2EE 1.3.

The Application Assembly Tool (AAT) has an option for migrating J2EE 1.2 application modules to J2EE 1.3. The J2EE 1.2 module is kept intact, with a new 1.3 module created. See also the `earconvert` tool documentation.

You must migrate J2EE 1.2 application modules to which you want to add J2EE 1.3 level WAR, RAR, EJB and client modules. This tool migrates only the application modules. J2EE 1.2-level modules inside a J2EE 1.2 application module must be migrated by other means.

3. Assemble new EJB modules (enterprise bean JAR files) as needed.

Assemble an EJB module to contain enterprise beans and related code artifacts. (Group Web components, client code, and resource adapter code in separate modules.)

You can install an EJB module as a standalone application or you can combine it with other modules into an enterprise application.

4. Assemble new Web modules (WAR files) as needed.

Assemble a Web module to contain servlets, JSP files, and related code artifacts. (Group enterprise beans, client code, and resource adapter code in separate modules.)

You can install a Web module as a standalone application or combine it with other modules into an enterprise application.

5. Assemble new application client modules (client JAR files) as needed.
6. Assemble new resource adapter archives (RAR files) as needed.

Assemble a resource adapter archive module to contain the library implementation code that your application uses to connect to enterprise information systems (EIS). (Group enterprise beans, Web components, and client code in separate modules.)

7. Assemble an application module from other module types.  
You are ready to combine your new or migrated modules into an application module (EAR file).  
For applications containing only Web modules, this step is optional. It is feasible to deploy Web modules without assembling them into application modules.
8. Verify your archive files.  
Verify your archive files and correct any problems so that generation of deployment code is successful. During verification, the Application Assembly Tool (AAT) checks that an archive file is complete, and that deployment descriptor properties and references contain appropriate values.
9. Remember to save your application one last time.
10. Generate code for deployment for applications containing EJB modules.  
If the application modules contain EJB modules, you must generate deployment code for the enterprise beans in the application before you deploy applications on the server. The Application Assembly Tool (AAT) provides this ability, or you can use the `ejbdeploy` command line tool.
11. Open existing modules (**File > Open**) in the AAT to modify them as needed.  
For example, you can add or remove modules and edit deployment descriptor properties.

#### What to do next

After assembling your applications, use a systems management tool to deploy the EAR or WAR files onto the application server.

The systems management tool follows the security and deployment instructions defined in the deployment descriptor, and enables you to modify bindings specified within the AAT. The tool locates the required external resources that the application uses, such as enterprise beans and databases.

Select a tool to use:

- Deploying and managing applications with the GUI
- Deploying and managing applications using programming (Java administrative APIs)
- Deploying and managing applications using scripting (WebSphere Application Server `wsadmin` tool)

If you are uncertain of which systems management tool to use, try using the administrative console.

---

## Application assembly and J2EE applications

Application assembly is the process of creating an Enterprise Archive (EAR) file containing all files related to an application, as well as an XML deployment descriptor for the application. This configuration and packaging prepares the application for deployment onto an application server.

EAR files are comprised of the following archives:

- Enterprise bean (JAR) files (known as EJB modules)
- Web application (WAR) files (known as Web modules)
- Application client (JAR) files (known as client modules)

- Resource adapter (RAR) files (known as resource adapter modules)

Ensure that modules are contained in an EAR file so that they may be deployed onto the server. The exceptions are WAR modules, which you can deploy individually. Although WAR modules can contain regular JAR files, they cannot contain the other module types described previously.

The assembly process includes the following:

- Selecting all of the files to include in the module
- Creating a deployment descriptor containing instructions for module deployment on the application server.

As you configure properties using the Application Assembly Tool (AAT), the tool generates the deployment descriptor for you. While the AAT graphical interface is recommended, you can also edit descriptors directly (`../javadoc/wccm/index.html`) in your favorite XML editor.

- Packaging modules into a single Enterprise archive (EAR) file, which contains one or more files in a compressed format

---

## Archive support in Version 5.0

These archives and Web components are supported:

- J2EE 1.3 Enterprise application (EAR) files
- EJB 2.0 (JAR) files
- Servlet 2.3 Web application WAR files
- Application Client 1.3 JAR files
- Connector 1.0 RAR files

These archive files and Web components are back-level and may be read but not created or changed:

- J2EE 1.2 EAR files
- EJB 1.1 JAR files
- Servlet 2.2 WAR files
- Application Client 1.2 JAR files

---

## Starting the Application Assembly Tool (AAT)

### Before you begin

A graphical interface is available for packaging code artifacts into various archives (modules) and configuring their J2EE 1.3 compliant deployment descriptors. The Application Assembly Tool (AAT) is available from the Windows Start menu, or you can invoke the tool from a command line as described in the Steps for this task.

If you access the Application Assembly Tool from a remote browser and select the Help, the Help files do not display. You can only view the Help files from a locally installed browser. To view the Help files and avoid this problem, close all the Netscape sessions on the remote machine and click **Help**. A new Netscape session starts, and you can then view the Help files.

### Steps for this task

1. Change directory at a system command prompt to the location of the assembly.bat|sh file, typically install\_root/bin.
2. Run the assembly script to launch the graphical interface.
3. Select whether to work with an existing module or create a new one.

### Results

The navigation tree displays a hierarchical structure used to build the contents of a new module, or to work with the contents of an existing module. Icons in the tree represent the components, assembly properties, and files for the module. The assembly properties appear in the AAT workspace.

Starting the Application Assembly Tool on a UNIX platform results in errors similar to the following:

```
...Font specified in font.properties not found [-urw-itc
zapfdingbats-medium-r-normal--*-%d--*-%p--*-%sun-fontspecific]
Font specified in font.properties not found [-urw-itc
zapfdingbats-medium-r-normal--*-%d--*-%p--*-%sun-fontspecific]
...
```

The Application Assembly Tool or installer functions are not affected by these errors. These messages display in the command shell that spawned the Java GUI. You can disregard these messages.

### What to do next

Consider whether you have any existing J2EE 1.2 application modules that you would like to migrate to J2EE 1.3.

You can create new modules of the following types, to assemble into an application module later:

- Assembling EJB modules
- Assembling Web modules
- Assembling application client modules
- Assembling resource adapter modules

Rather than create new modules to assemble an application, you can proceed directly to assembling a new application module. While assembling an application module, you can create any new modules that you need.

---

## Migrating application modules from J2EE 1.2 to J2EE 1.3

The Application Assembly Tool (AAT) has an option for migrating J2EE 1.2 application modules to J2EE 1.3. The J2EE 1.2 module is kept intact, with a new 1.3 module created. See also the earconvert tool documentation.

Migrate J2EE 1.2 application modules to which you want to add J2EE 1.3 level Web application (WAR) modules, Resource adapter (RAR) modules, Entity bean (EJB) modules, and application client modules. This tool migrates only the application modules. Migrate J2EE 1.2-level modules inside a J2EE 1.2 application module by other means.

**Note:** When Entity beans are moved from a J2EE 1.2 module to a J2EE 1.3 module, the EJB container will then apply rules defined in the EJB 2.0 specification to these

beans. The EJB 2.0 specification mandates that when a `findBy` method is called on a bean home (except for `findByPrimaryKey`), the EJB container must cause other Entity beans enlisted in the same transaction to write out their current state to the persistent store. This is to ensure that the `findBy` operation is performed on the most current data. Application developers should plan for and be aware of any changes to the application behavior as a result of this rule.

#### Steps for this task

1. Start the AAT.
2. Use it to open the J2EE 1.2 application module you want to migrate.
3. Click **Convert EAR** from the file menu.
4. Save the new J2EE 1.3 application.

#### What to do next

Assemble zero or more new modules of your choice:

- Assembling EJB modules
- Assembling Web modules
- Assembling application client modules
- Assembling resource adapter modules

Another option is to proceed directly to assembling a new application module. You can create any new modules that you need, while assembling an application module.

## earconvert tool

A command line tool is provided for migrating J2EE 1.2 application modules to J2EE 1.3. This migration enables you to add J2EE 1.3 modules to the migrated application module. See also the Application Assembly Tool (AAT) for information on performing this task.

Migrate J2EE 1.2 application modules to which you want to add J2EE 1.3 level Web application (WAR), Resource adapter (RAR), Enterprise beans (EJB), and client modules. This tool migrates only the application modules. Ensure that you migrate J2EE 1.2-level modules inside a J2EE 1.2 application to prevent working with back-level files.

### Syntax

Locate the file that contains the `earconvert.bat|sh` tool, which must reside in the `bin` directory of the product installation root.

#### **earconvert**

*j2ee\_1.2\_file\_name*

*j2ee\_1.3\_file\_name*

### Parameters

Supported arguments include:

*"j2ee\_1.2\_file\_name"*

Specifies the actual name of the existing J2EE 1.2 application file. (In this and other arguments, use quotation marks to allow for path names that contain spaces.)

*"j2ee\_1.3\_file\_name"*

Specifies what you would like to name the new J2EE 1.3 application file.

### Examples

The following command creates a new J2EE 1.3 archive, `new_application.ear`, based on the J2EE 1.2 archive, `existing_application.ear`.

```
earconvert existing_application.ear new_application.ear
```

---

## Assembling new or modifying existing modules

### Before you begin

Ensure that code artifacts, such as servlets, JSP files, enterprise beans, and application clients are assembled into their respective modules.

If you want to use existing J2EE 1.2 modules in your J2EE 1.3 application, migrate these modules to J2EE 1.3 first. Also migrate any J2EE 1.2 application modules to which you want to add J2EE 1.3 modules.

You are now ready to combine your new or migrated modules into an application module Enterprise application (EAR file).

The Application Assembly Tool (AAT) provides flexibility in assembling applications from various Web application (WAR), Resource adapter (RAR), Enterprise beans (EJB JAR), and application client (JAR) files. Options described in assembling applications include:

- Importing an existing module (JAR, RAR or WAR file)
- Creating a new module while you create the new application
- Copying code artifacts, such as servlets, from one module to another of the same type, to reside in the new application

### Steps for this task

1. Start the AAT.
2. From the New tab, select **Application**, and click **OK**, if you did not already specify to create a new application module.  
Each of the next three steps is optional, but you must perform at least one of them.
3. **(Optional)** Import existing modules into the application module.
  - a. Right-click the folder for the type of module you want to import, such as an EJB module, in the navigation tree.
  - b. Click **Import** from its right-click menu.
  - c. Use the file browser to locate and select the archive file for the module.
  - d. Click **Open**. The archive file appears under the appropriate folder in the navigation tree.
  - e. Click the plus sign (+) next to the icon for the archive, to view the module contents and edit its properties if needed.
  - f. Save the application module.
4. **(Optional)** Create a new archive file to include in the application.



- a. Right-click the folder for the type of module to create (such as enterprise beans (EJB) modules, Web application modules (WAR), resource adapter (RAR) files, or application client modules) in the navigation tree.
  - b. Click **New** from its right-click menu.
  - c. Configure properties of the new module when it displays.
  - d. Click **OK**. The archive file displays under the appropriate folder.
  - e. Click the plus sign (+) to verify file contents and enter assembly properties.
  - f. Add enterprise beans, if this is an EJB module.
  - g. Right-click the folder corresponding to the type of bean to create (session bean or entity bean), and click **New** or **Import**.
  - h. Configure properties of the enterprise bean when it displays.
  - i. Click **OK**. The enterprise bean appears in the navigation pane.
  - j. Click the plus sign (+) to verify file contents and enter assembly properties.
  - k. Save the application module.
5. **(Optional)** Copy code artifacts, such as servlets, from one module to another of the same type, to reside in the new application.
    - a. Identify the code artifact to copy, and the type of module in which it resides. Make sure you already have the same kind of module (such as a Web module) created in the new application module.
    - b. Open a separate, existing module in the AAT by selecting **File > Open** from the menu bar.
    - c. Arrange the AAT workspace so that you can see both the new application module and the source archive containing the code artifact.
    - d. Copy and paste the code artifact from the source module to the same module type in the new application.
 

For example, copy a container-managed persistence (CMP) bean from the source EJB module into the new EJB application module.
    - e. Save the application module.
  6. Continue to add desired modules to the application module.
  7. Define security properties for the application.
    - a. Right-click the **Security Roles** icon in the navigation tree.
    - b. Click **New**.
    - c. Configure the security properties.
    - d. Click **OK**.
  8. Add supplementary files needed by the application.
    - a. Right-click the **Files** icon in the navigation tree, and select **Add Files**.
    - b. Add files, using the Add Files dialog.
  9. Save the application module.

### Results

You are performing application assembly results in a J2EE 1.3 compliant EAR file containing one or more WAR, RAR, or JAR files.

**Note:** If you use the Application Assembly Tool to create application client modules, you must also use the Application Client Resource Configuration Tool. Using this tool, you can define references to resources (other than enterprise beans) on the machine where the application client resides.

## What to do next

"Verifying archive files". "Generating code for deployment".

After an application is assembled and you generate your application code for deployment, use the administrative console to install the application in an application server.

Use the administrative console at installation time to carry out the security instructions defined in the deployment descriptor and to locate required external resources, such as enterprise beans and databases. You can add configuration properties and redefine binding properties defined in the Application Assembly Tool.

After the application deploys, you can use the Application Assembly Tool to modify the application by adding or removing modules, editing deployment descriptor properties and regenerating code for deployment.

**Note:** If your application has a large number of modules, it might not install successfully onto a server. Package your application so the .ear file has as few modules as are necessary. Modules can include metadata for the modules such as information on deployment descriptors, bindings and IBM extensions.

## **Adding files to assembled modules**

### Before you begin

Review the usage scenario (as follows) to become familiar with the **Add Files** dialog.

Use the **Add Files** dialog box of the Application Assembly Tool (AAT) to import files into assembled modules including Enterprise application (EAR), Web application (WAR), Resource adapter (RAR) and Application client (JAR) files.

This task assumes that you are performing another task, such as assembling a Web or EJB module, when the **Add Files** dialog is presented to you.

### Steps for this task

1. Click **Browse**. Locate the files to add.
  - To add specific individual files, select the directory or archive (WAR, JAR, RAR, ZIP, for example) containing the files.
  - To add an entire directory of files, select its parent directory.
2. Click **OK**.


The selected directory or archive appears in the top left part of the dialog box, in an expandable tree.

The top right part of the dialog box shows the contents of the directory, subdirectory, or archive that is selected on the left-hand side.
3. Select one or more items to add from the top right part of the dialog, then click **Add**.

As you add files, they will be displayed in the lower half of the dialog box.
4. Change your left-hand selection to gain access to other files that you want to add, as needed.
5. Click **OK** when all of the files that you want to add appear in the lower half of the dialog box.

## Usage scenario

The following example refers to the main areas of the Add Files dialog box. Details such as clicking OK are omitted. Refer to the detailed task steps above for this information.



Suppose you are constructing a new application module and want to add `myFile.txt` file to the archive as a supplementary file. The `myFile.txt` currently is contained within the `myFiles` subdirectory of a JAR file that resides somewhere on your directory system.

1. Browse for the JAR file.
2. Select the JAR file.
3. Exit the browse dialog. At this point:
  - Area 1 of the **Add Files** dialog contains the path to the JAR file.
  - Area 2 displays the JAR file name as the root directory of an expandable tree showing the directories in the JAR file — including `myFiles`.
  - Area 3 shows the root contents of the JAR file, as well as any subdirectories visible from the root.
4. Select the `myFiles` directory from area 2, causing `myFile.txt` to become visible in area 3.
5. Click **myFiles.txt** from area 3 and specify to **Add** the file. Now this file is listed in area 4, the lower half of the dialog, which indicates it is the file that you want to add to the new application module.
6. Exit the **Add Files** dialog.

## Resource environment reference assembly settings

Resource environment reference elements contain declarations of an enterprise bean's reference to an administered object associated with a resource in the enterprise bean's environment.

### Name

Specifies the name of the resource environment reference.

Its value is the environment entry name used in the enterprise bean code.

Data type

String

**Description**

Contains the information that the EJB jar file producer wants to provide to the EJB jar file consumer.

Data type String

**Type**

Specifies the type of a resource environment reference.

Data type String

## Resource Adapter Archive file assembly settings

Use this page to set the resource adapter archive file properties.

**File name**

Specifies the file name of the Resource Adapter Archive.

Data type String

**Display name**

Specifies a short name that is intended to be displayed by the GUI.

Data type String

**Description**

Specifies a description that should include any information that the component file producer wants to provide to the consumer of the component file (that is, to the deployer).

Data type String

**EIS type**

This helps in identifying EIS instances that can be used with this resource adapter.

Data type String

**Vendor name**

Specifies a string-based version of the resource adapter from the resource adapter provider.

Data type String

**Version**

Specifies a string-based version of the resource adapter from the resource adapter provider.

Data type String

**Specification**

Specifies the version of the connector architecture specification that is supported by this resource adapter.

Data type String

### License required

Specifies if a license is or is not required.

#### Description

If a license is required, this field specifies the licensing requirements for the resource adapter module. For example, duration of license, number of connection restrictions, and so forth.

Data type String

### Implementation

The element (managedconnectionfactory class) that specifies the fully qualified name of the Java class that implements the javax.resource.spi.ManagedConnectionFactory interface.

Data type Class

### Interface

The element (credential-interface) that specifies the interface that the resource adapter implementation supports for the representation of the credentials.

The possible values are:

```
<credential-interface> javax.resource.spi.security.PasswordCredential
</credential-interface>
<credential-interface> javax.resource.spi.security.GenericCredential
</credential-interface>
```

Data type Class

### Implementation

The element (connectionfactory class) that specifies the fully-qualified name of the ConnectionFactory class that implements the resource adapter specific ConnectionFactory interface.

Data type Class

### Interface

The element (connection-interface) that specifies the fully-qualified name of the Connection interface supported by the resource adapter.

Data type Class

### Implementation

The element (connection class) that specifies the fully-qualified name of the Connection class that implements the resource adapter specific Connection interface.

Data type Class

### Support Reauthentication

Specifies whether the resource adapter implementation supports re-authentication of existing ManagedConnection instances.

The values are either True or False.

Data type       String

### **Transaction**

Specifies the level of transaction support provided by the resource adapter.

The three possible values are:

- NoTransaction
- LocalTransaction
- XATransaction

Data type       String

### **Small Icon**

The image is used as an icon to represent the module in a GUI.

Specifies a JPEG or GIF file containing a small image (16x16 pixels).

Data type       Image

### **Large Icon**

The image is used as an icon to represent the module in a GUI.

Specifies a JPEG or GIF file containing a small image (32x32 pixels).

Data type       Image

### **Basic Password**

The basic user password authentication mechanism that is specific to an EIS.

#### **Credential Interface**

Specifies the interface that the resource adapter implementation supports for the representation of the credentials. For Basic Password the credential value is *javax.resource.spi.security.PasswordCredential*.

#### **Description**

Any information that describes Basic Password selection.

Data type       String

### **Kerberos V5**

Specifies a Kerberos version 5 authentication mechanism.

#### **Credential Interface**

Specifies the interface that the resource adapter implementation supports for the representation of the credentials. For Kerberos version 5, the credential value is *javax.resource.spi.security.GenericCredential*.

#### **Description**

Any information that describes the Kerberos V5 selection.

Data type       String

### Property Name

Specifies the name of a configuration property.

The possible values are:

```
<config-property-name>ServerName</config-property-name>
<config-property-name>PortNumber</config-property-name>
<config-property-name>UserName</config-property-name>
<config-property-name>Password</config-property-name>
<config-property-name>ConnectionURL</config-property-name>
```

Data type          String

### Property Type

Contains the fully-qualified Java type of a configuration property as required by the ManagedConnectionFactory instance.

Data type          String

### Property Value

Contains the value of a configuration entry.

Data type          String

### Description

Describes the parent element.

Data type          String

### Permission Specification

Specifies a security permission that is required by the resource adapter code.

Data type          String

---

## Saving applications after assembly

Periodically save modules that you assemble with the Application Assembly Tool (AAT). Save any changes right before you close the module with which you are working.

### Before you begin

This task assumes you have started the AAT and are working with a particular module.

### Steps for this task

1. Save the archive file by clicking **File > Save As**.
  - If you are saving an existing archive file or application, click **File > Save**.
2. Name the new archive file or application whatever you like.

This step is optional if you are working with an existing archive file or application.

### What to do next

Now that you have saved your assembled application, you can verify your archives and generate code for deployment.

---

## Verifying archive files

Verify your archive files and correct any problems so that generation of deployment code is successful. During verification, the Application Assembly Tool (AAT) checks that an archive file is complete, and that deployment descriptor properties and references contain appropriate values.

This task assumes you have previously assembled and saved one or more modules.

### Steps for this task

1. Start the Application Assembly Tool (AAT).
2. Click **File > Open** and select the module to verify.
3. Right-click the name of the module at the top of the navigation pane and click **Verify**.
4. Click **Verify** in the **Verify** window. The tool displays a scrolling window for viewing status messages as the verification proceeds.
5. Save the application.

### Results

Archive files have been verified. The following list includes, but is not limited to, areas that the verification process has checked:

- Required deployment properties contain values.
- Values specified for environment entries match their associated Java types.
- In both Enterprise application (EAR) and Web application (WAR) files:
  - The target enterprise bean of the link exists for EJB references.
  - The target role exists for security role references.
  - Security roles are unique.
- Each module listed in the deployment descriptor exists in the archive for EAR files.
- Files for icons, servlets, error and welcome pages listed in the deployment descriptor have corresponding files in the archive for WAR files.
- For EJB modules:
  - All class files referenced in the deployment descriptor exist in the JAR file.
  - Method signatures for enterprise bean home, remote and implementation classes are compliant with the EJB 2.0 specification.

### What to do next

If your application module contains EJB modules, generate code for deployment.

Otherwise, you are ready to deploy this application module (or standalone Web module) onto the application server.



---

## Application assembly performance checklist

Application assembly tools are used to build J2EE components and modules into J2EE applications. Generally, assembling consists of defining application components and their attributes including enterprise beans, servlets and resource references. Many of these application configuration settings and attributes play an important role in the run-time performance of the deployed application. Use the following information as a check list of important parameters and advice for finding optimal settings:

- EJB modules
  - Entity bean Bean Cache - Activate at and Bean Cache - Load at settings
  - Method extensions Isolation level and Access intent settings
  - (Container transactions assembly settings)
- Web modules
  - Web modules assembly settings
    - Distributable
    - Reload interval
    - Reload enabled
- Web components
  - Load on startup

---

## Generating code for deployment

Before deploying applications on the server, if the application modules contain EJB modules, you must generate deployment code for the enterprise beans in the application. The Application Assembly Tool (AAT) provides this ability, or you can use the `ejbdeploy` command line tool.

### Before you begin

This task assumes you have already assembled an EJB module, added it to an application module, saved the application module, and verified the application module.

Before installing your application in WebSphere Application Server, you must generate deployment code for the application. This step is required for EJB modules and for any Enterprise application (EAR) files that contain EJB modules. During code generation, the Application Assembly Tool invokes the `EJBDeploy` tool to prepare entity bean (JAR) files for deployment in run time environment. To deploy a J2EE application, you can install the application in the administrative console.

The following steps assume that you are using the Application Assembly Tool to generate code for deployment.

### Steps for this task

1. Start the Application Assembly Tool (AAT).
2. Open the EAR or JAR file for which you want to generate code for deployment.
3. Click **File > Generate code for deployment** from the menu bar.
4. Specify the options for the server to use for generating code for the application deployment.

**Note:** For Container managed persistence (CMP) entity beans, if the JAR file that you opened (inputJar file) contains a map and schema document, that schema is used. If the JAR file does not contain a map and schema document, the Application Assembly Tool uses a top-down mapping to generate files that contain mapping and database schema information.

5. Click **Generate Now**.

Review the messaging box for details of any error that might occur.

**Note:** Do not change the default output file name to be the same as the input filename, as the AAT cannot read and write to the same file name, and therefore, an error will occur.

### Results

After deployment code is generated for an application, the deployable archive is renamed with the prefix Deployed\_.

### What to do next

Install the application on your server machine.

**Note:** Before deploying the application in your run time environment, you might need to set classpaths.

## **ejbdeploy tool**

You can generate code for deployment by either using the Application Assembly Tool (AAT) or by using the Deployment Tool for Enterprise Java Beans (ejbdeploy) from a command prompt. For example, the options that you are able to set in AAT correspond with commands that the EJBDeploy tool uses to generate code for deploying an application.

### **Syntax**

For a detailed list of available options in the EJBDeploy tool, enter `ejbdeploy` from a command prompt.

## **ejbdeploy syntax — relationship to Application Assembly Tool options**

### **Abstract:**

Many options for generating code for deployment in the Application Assembly Tool (AAT) directly correlate with the options in the command line EJBDeploy tool.

| <b>Application Assembly Tool options</b> | <b>EJBDeploy tool options</b> |
|--|-------------------------------|
| Deployed module location                 | <code>outputJar</code>        |
| Working Directory                        | <code>workingDirectory</code> |
| Dependent classpath                      | <code>cp</code>               |
| Code generation only                     | <code>codegen</code>          |
| Verify archive (unchecked)               | <code>novalidate</code>       |
| RMIC options                             | <code>rmic options</code>     |

|               |          |
|---------------|----------|
| Database type | dbvendor |
| Database name | dbname   |
| Schema name   | dbschema |

---

## Application Assembly Tool: Resources for learning




Use the following links to find relevant supplemental information about the Application Assembly Tool. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.



View links to additional information about:

- Programming specifications
- Administration

### Programming specifications

-  **J2EE 1.3 specification** ([http://java.sun.com/j2ee/j2ee-1\\_3-fr-spec.pdf](http://java.sun.com/j2ee/j2ee-1_3-fr-spec.pdf))
-  **EJB specifications** (<http://java.sun.com/products/ejb/docs.html> )
-  **Servlet specifications**  
(<http://java.sun.com/products/servlet/download.html>)

### Administration

-  **Application Client files**  
(<http://developer.java.sun.com/developer/technicalArticles/J2EE/appclient/>)
-  **Connector RAR files** (<http://java.sun.com/j2ee/connector/>)



---

## Chapter 28. Deploying and managing applications

After you develop an enterprise application and configure an application server, you can use the administrative console to install application files on the server and manage the activity of deployed applications.

### Steps for this task

1. Install your application on your application server.
2. Start and stop applications.
3. Edit the administrative configuration for an application. Go to the settings page for an application, change the values for settings as needed, and click **OK**.
4. **(Optional)** Export applications.
5. **(Optional)** Export DDL files.
6. Update application binary files.
7. Uninstall applications.

After making changes to administrative configurations of your applications, ensure that you click **Save** on the administrative console taskbar to save the changes.

---

### Enterprise applications

Enterprise applications (or J2EE applications) are applications that conform to the Java 2 Platform, Enterprise Edition, specification.

Enterprise applications can consist of the following:

- Zero or more EJB modules
- Zero or more Web modules
- Zero or more connector modules (packaged in RAR files)
- Zero or more application client modules
- Optionally, additional JAR files containing dependent classes or other components required by the application
- Any combination of the above

A J2EE application is represented by, and packaged in, an enterprise archive (EAR) file.

---

### Installing a new application

To install an enterprise application to a WebSphere Application Server configuration, you can use the administrative console or the wsadmin tool. The steps below describe how to use the administrative console to install an application, EJB component, or Web module.

**Note:** Once you start performing the steps below, click **Cancel** to exit if you decide not to install the application. Do not simply move to another administrative console page without first clicking **Cancel** on an application installation page.

### Steps for this task

1. Click **Applications > Install New Application** in the console navigation tree. The first of two Preparing for application install pages is shown.
2. On the first Preparing for application install page:
  - a. Specify the full path name of the source application file (.ear file otherwise known as an EAR file). The EAR file that you are installing can be either on the client machine (the machine that runs the Web browser) or on the server machine (the machine to which the client is connected). If you specify an EAR file on the client machine, then the administrative console uploads the EAR file to the machine on which the console is running and proceeds with application installation. You can also specify a stand-alone WAR or JAR file for installation.
  - b. If you are installing a stand-alone WAR file, specify the context root.
  - c. Click **Next**.
3. On the second Preparing for application install page:
  - a. Select whether to generate default bindings. Using the default bindings causes any incomplete bindings in the application to be filled in with default values. Existing bindings are not altered. You can customize default values used in generating default bindings. For example, you can specify JNDI prefix for all the EJB files in EJB modules, default data source and connection factory settings for EJB modules, virtual host for web modules, and so on. "Preparing for application installation settings" describes available customizations and provides sample bindings.
  - b. Click **Next**. The Install New Application pages are now shown. If you chose to generate default bindings, you can proceed to the Summary step (step 23 below). "Example: Installing an EAR file using the default bindings" provides sample steps.
4. On the **Step: Provide options to perform the installation** panel, provide values for the following settings specific to WebSphere Application Server. Default values are used if you do not specify a value.
  - a. For **Pre-compile JSP**, specify whether to precompile JSP files as a part of installation. The default is not to precompile JSP files.
  - b. For **Directory to Install Application**, specify the directory to which the application EAR file will be installed. The default value is the value of `APP_INSTALL_ROOT/cell_name`, where the `APP_INSTALL_ROOT` variable is `install_root/installedApps`; for example, `C:\WebSphere\AppServer\installedApps\cell_name`.

**Note:** If an installation directory is not specified when an application is installed on a single-server (base) configuration, the application is installed in `APP_INSTALL_ROOT/base_cell_name`. When the base server is made a part of a Network Deployment configuration (using the `addNode` utility), the cell name of the new configuration becomes the cell name of the deployment manager node. If the `-includeapps` option is used for the `addNode` utility, then the applications that are installed prior to the `addNode` operation still use the installation directory `APP_INSTALL_ROOT/base_cell_name`. However, an application that is installed after the base server is added to the network configuration uses the default installation directory `APP_INSTALL_ROOT/network_cell_name`. To move the application to the `APP_INSTALL_ROOT/network_cell_name` location upon running the `addNode` operation, you should explicitly specify the installation directory as `${APP_INSTALL_ROOT}/${CELL}` during installation. In such a case, the application files can always be found under `APP_INSTALL_ROOT/current_cell_name`.

- c. For **Distribute Application**, specify whether WebSphere Application Server expands or deletes application binaries in the installation destination. The default is to enable application distribution. As a result, when you save changes in the console, application binaries for newly installed applications are expanded to the directory specified. The binaries are also deleted when you uninstall and save changes to the configuration. If you disable this option, then you must ensure that the application binaries are expanded appropriately in the destination directories of all nodes where the application is expected to run.
  - d. For **Use Binary Configuration**, specify whether the application server uses the binding, extensions, and deployment descriptors located with the application deployment document, the deployment.xml file (default), or those located in the EAR file. The default is not to use the binary configuration.
  - e. For **Deploy EJBs**, specify whether the EJBDeploy tool runs during application installation. The tool generates code needed to run EJB files. The default is not to run the EJBDeploy tool. You must enable this setting if the EAR file was assembled using the Application Assembly (AAT) tool and the EJBDeploy tool was not run during assembly, if the EAR file was not assembled using the AAT tool, or if the EAR file was assembled using versions of the AAT tool previous to Version 5. Note that enabling this setting might cause the installation program to run for several minutes.
  - f. For **Application Name**, name the application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names.
  - g. For **Create MBeans for Resources**, specify whether to create MBeans for various resources (such as servlets or JSP files) within an application when the application is started. The default is to create MBean instances.
  - h. For **Enable class reloading**, specify whether to enable class reloading when application files are updated. The default is not to enable class reloading.
  - i. For **Reload Interval**, specify the number of seconds to scan the application's file system for updated files. The default is the value of the reload interval attribute in the IBM extension (META-INF/ibm-application-ext.xmi) file of the EAR file. This setting takes effect only if class reloading is enabled.  
  
The reload interval specified here overrides the value specified in the IBM extensions for each Web module in the EAR file (which in turn overrides the reload interval specified in the IBM extensions for the application in the EAR file).
5. If your application uses EJB modules, on the **Step: Provide JNDI Names for Beans** panel, specify a JNDI name for each enterprise bean in every EJB module. You must specify a JNDI name for every enterprise bean defined in the application. For example, for the EJB module MyBean.jar, specify MyBean.
  6. If your application uses EJB modules that contain Container Managed Persistence (CMP) beans that are based on the EJB 1.x specification, for **Step: Provide default datasource mapping for modules containing 1.x entity beans**, specify a JNDI name for the default data source for the EJB modules. The default data source for the EJB modules is optional if data sources are specified for individual CMP beans.
  7. If your application has CMP beans that are based on the EJB 1.x specification, for **Step: Map datasources for all 1.x CMP**, specify a JNDI name for data sources to be used for each of the 1.x CMP beans. The data source attribute is

optional for individual CMP beans if a default data source is specified for the EJB module that contains CMP beans. If neither a default data source for the EJB module nor a data source for individual CMP beans are specified, then a validation error displays after you click **Finish** (step 23) and the installation is cancelled.

8. If your application defines EJB references, for **Step: Map EJB references to beans**, specify JNDI names for enterprise beans that represent the logical names specified in EJB references. Each EJB reference defined in the application must be bound to an EJB file before clicking **Finish** on the Summary panel.
9. If your application defines resource references, for **Step: Map resource references to resources**, specify JNDI names for the resources that represent the logical names defined in resource references. Each resource reference defined in the application must be bound to a resource defined in your WebSphere Application Server configuration before clicking on **Finish** on the Summary panel.
10. If your application uses Web modules, for **Step: Map virtual hosts for web modules**, select a virtual host from the list that should map to a Web module defined in the application. The port number specified in the virtual host definition is used in the URL that is used to access artifacts such as servlets and JSP files in the Web module. Each Web module must have a virtual host to which it maps. Not specifying all needed virtual hosts will result in a validation error displaying after you click **Finish** on the Summary panel.
11. On the **Step: Map modules to application servers** panel, for every module select a target server or a cluster from the **Clusters and Servers** list. Place a checkmark in the check box beside **Module** to select all of the application modules or select individual modules.
12. If the application has security roles defined in its deployment descriptor then, for **Step: Map security roles to users/groups**, specify users and groups that are mapped to each of the security roles. Place a checkmark in the check box beside **Role** to select all of the roles or select individual roles. For each role, you can specify if predefined users such as **Everyone** or **All Authenticated users** are mapped to it. To select specific users or groups from the user registry:
  - a. Select a role and click **Lookup users** or **Lookup groups**.
  - b. On the Lookup users/groups panel shown, enter search criteria to extract a list of users or groups from the user registry.
  - c. Select individual users or groups from the results displayed.
  - d. Click **OK** to map the selected users or groups to the role selected on the **Step: Map security roles to users/groups** panel.
13. If the application has RunAs roles defined in its deployment descriptor, for **Step: Map RunAs roles to user**, specify the RunAs user name and password for every RunAs role. RunAs roles are used by enterprise beans that must run as a particular role while interacting with another enterprise bean. Place a checkmark in the check box beside **Role** to select all of the roles or select individual roles. After selecting a role, enter values for the user name, password, and verify password and click **Apply**.
14. If your application contains EJB 1.x CMP beans that do not have method permissions defined for some of the EJB methods, for **Step: Ensure all unprotected 1.x methods have the correct level of protection**, specify if you want to leave such methods unprotected or assign protection with deny all access.



15. If your application contains message driven enterprise beans, for **Step: Provide Listener Ports for messaging beans**, provide a listener port name for every message driven bean. If a name is not specified for each bean, then a validation error displays after you click on **Finish** on the Summary panel.
16. If your application uses EJB modules that contain CMP beans that are based on the EJB 2.0 specification, for **Step: Provide default datasource mapping for modules containing 2.0 entity beans**, specify a JNDI name for the default data source and the type of resource authorization to be used for the default data source for the EJB modules. The default data source for EJB modules is optional if data sources are specified for individual CMP beans.
17. If your application has CMP beans that are based on the EJB 2.0 specification, on the **Step: Map datasources for all 2.0 CMP** panel, for each of the 2.0 CMP beans specify a JNDI name and the type of resource authorization for data sources to be used. The data source attribute is optional for individual CMP beans if a default data source is specified for the EJB module that contains CMP beans. If neither a default data source for the EJB module nor a data source for individual CMP beans are specified, then a validation error is shown after you click **Finish** and installation is cancelled.
18. If your application contains EJB 2.0 CMP beans that do not have method permissions defined in the deployment descriptors for some of the EJB methods, on the **Step: Ensure all unprotected 2.0 methods have the correct level of protection** panel, specify whether you want to assign a specific role to the unprotected methods, add the methods to the exclude list, or mark them as unchecked. Methods added to the exclude list are marked as uncallable. For methods marked unchecked no authorization check is performed prior to their invocation.
19. If the **Deploy EJBs** setting is enabled on the **Provide options to perform the installation** panel, then you can specify options for the EJBDeploy tool on the **Step: Provide options to perform the EJB Deploy** panel. On this panel, you can specify extra classpath, rmic options, database types, and database schema names to be used while running the EJBDeploy tool. The tool is run on the EAR file during installation after you click **Finish**.
20. If your application contains resource environment references, for **Step: Mapping Resource Environment References to Resources**, specify JNDI names of resources that map to the logical names defined in resource environment references. If each resource environment reference does not have a resource associated with it, a validation error is shown after you click **Finish**.
21. If your application defines **Run-As Identity** as *System Identity*, for **Step: Replacing RunAs System to RunAs Roles**, you can optionally change it to *Run-As role* and specify a user name and password for the RunAs role specified. Selecting *System Identity* implies that the invocation is done using the WebSphere Application Server security server ID and should be used with caution as this ID has more privileges.
22. If your application has resource references that map to resources that have an Oracle database doing back-end processing, for **Step: Specify the isolation level for Oracle type provider**, specify or correct the isolation level to be used for such resources when used by the application. Oracle databases support ReadCommitted and Serializable isolation levels only.
23. On the Summary panel, verify the cell, node, and server onto which the application modules will install. Beside the **Cell/Node/Server** option, click **Click here** and verify the settings. Then click **Finish**.

**Note:** After clicking **Finish**, if you receive an OutOfMemory exception and the source application file does not install, your system might not have enough

memory or your application might have too many modules in it to install successfully onto the server. If lack of system memory is not the cause of the exception, package your application again so the .ear file has fewer modules. If lack of system memory and the number of modules are not the cause of the exception, check the options you specified on the Java Virtual Machine page of the application server running the administrative console. Then, try installing the application file again.

24. Associate any shared libraries that the application needs to the application.
25. Click **Save** on the administrative console taskbar to save the changes to your configuration. The application is registered with the administrative configuration and application files are copied to the target directory, which is *install\_root/installedApps/cell\_name* by default or the directory that you designate. For the single-server (base) installation, application files are copied to the destination directory when you click **Save**; for the Network Deployment installation, files are copied to remote nodes when the configuration on the deployment manager synchronizes with the configuration on individual nodes.
26. Test the application. For example, point a Web browser at the URL for the deployed application and examine the performance of the application. If necessary, update the application.

## Preparing for application installation settings

Use this page to install an application (EAR file) or module (JAR or WAR file).

To view this administrative console page, click **Applications > Install New Application**.

Follow the steps on this page to install an application or module. You must complete, at minimum, the first step; you must complete some or all of the later steps, depending on whether you are installing an application, EJB module, or Web module.

### Path

Specifies the fully qualified path to the .ear, .jar, or .war file for the enterprise application.

Use **Local path** if the browser and application files are on the same machine (whether or not the server is on that machine, too).

Use **Server path** if the application file resides on any node in the current cell context. You can browse the entire file system of a node if the node agent or deployment manager is running on that selected node. Only .ear, .jar, or .war files are shown during the browsing.

During application installation, application files are typically uploaded from a client machine running the browser to the server machine running the administrative console, where they are deployed. In such cases, the Web browser running the administrative console is used to select EAR, WAR, or JAR modules to upload to the server machine.

In some cases, however, the application files reside on the file system of any of the nodes in a cell. To have the application server install these files, use the **Server path** option.

You can also use this option to specify an application file already residing on the machine running the application server. For example, the field value on Windows NT might be C:\WebSphere\AppServer\installableApps\test.ear. If you are installing a stand-alone WAR module, then you also must specify the context root.

### Context Root

Specifies the context root of the Web application (WAR).

This field is used only to install a stand-alone WAR file. The context root is combined with the defined servlet mapping (from the WAR file) to compose the full URL that users type to access the servlet. For example, if the context root is /gettingstarted and the servlet mapping is MySession, then the URL is http://host:port/gettingstarted/MySession.

### Generate Default Bindings

Specifies whether to generate default bindings. If you place a checkmark in the check box, then any incomplete bindings in the application are filled in with default values. Existing bindings are not altered.

By choosing this option, you can directly jump to the Summary step and install the application if none of the steps have a red asterisk (\*) next to them. A red asterisk denotes that the step has incomplete data and requires a valid value. On the Summary panel, verify the cell, node and server on which the application is installed.

Bindings are generated as follows:

- EJB JNDI names are generated of the form *prefix/ejb-name*. The default prefix is *ejb*, but can be overridden. The *ejb-name* is as specified in the deployment descriptors `<ejb-name>` tag.
- EJB references are bound as follows: If an `<ejb-link>` is found, it is honored. Otherwise, if a unique enterprise bean is found with a matching home (or local home) interface as the referenced bean, the reference is resolved automatically.
- Resource reference bindings are derived from the `<res-ref-name>` tag. Note that this action assumes that the `java:comp/env` name is the same as the resource global JNDI name.
- Connection factory bindings (for EJB 2.0 JAR files) are generated based on the JNDI name and authorization information provided. This action results in default connection factory settings for each EJB 2.0 JAR file in the application being installed. No bean-level connection factory bindings are generated.
- Data source bindings (for EJB 1.1 JAR files) are generated based on the JNDI name, data source user name password options. This results in default data source settings for each EJB JAR file. No bean-level data source bindings are generated.
- Message Driven Bean (MDB) listener ports are derived from the MDB `<ejb-name>` tag with the string `Port` appended.
- For `.war` files, the virtual host is set as `default_host` unless otherwise specified.

The default strategy suffices for most applications or at least for most bindings in most applications. However, it does not work if:

- You want to explicitly control the global JNDI names of one or more EJB files.
- You need tighter control of data source bindings for CMPs. That is, you have multiple data sources and need more than one global data source.
- You must map resource references to global resource JNDI names that are different from the `java:comp/env` name.

In such cases, you can alter the behavior with an XML document (a custom strategy). Use the **Specific bindings file** field to specify a custom strategy and see the field's help for examples.

### Prefixes

Specifies prefixes to use for generated JNDI names.

### Override

Specifies whether to override existing bindings.

If this check box is checked, the existing bindings are overridden by the generated ones.

### EJB 1.1 CMP bindings

Specifies the default data source JNDI name.

If the **Default Bindings for EJB 1.1 CMPs** radio button is selected, specify the JNDI name for the default data source to be used with the CMP 1.1 beans. Also specify the user ID and password for this default data source.

### Connection Factory Bindings

Specifies the default data source JNDI name.

If the **Default connection factory bindings** radio button is selected, specify the JNDI name for the default data source to be used with the bindings. Also specify the resource authorization.

### Virtual Host

Specifies the virtual host for WAR modules.

### Specific bindings file

Specifies a bindings file that overrides the default binding.

Alter the behavior of the default binding with an XML document (*aka* custom strategy). Custom strategies extend the default strategy so you only need to customize those areas where the default strategy is insufficient. That is, you only need to describe how you want to change the bindings generated by the default strategy; you do not have to define bindings for the entire application.

Brief examples of how to override various aspects of the default bindings generator follow:

#### Controlling an EJB JNDI name

```
<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
    <ejb-jar-binding>
      <jar-name>helloEjb.jar</jar-name>
      <!-- this name must match the module name in the .ear file -->
      <ejb-bindings>
        <ejb-binding>
          <ejb-name>HelloEjb</ejb-name>
          <!-- this must match the <ejb-name> entry in the EJB jar DD -->
          <jndi-name>com/acme/ejb/HelloHome</jndi-name>
        </ejb-binding>
      </ejb-bindings>
    </ejb-jar-binding>
  </module-bindings>
</dfltbndngs>
```

### Setting the connection factory binding for an EJB JAR file

```
<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
    <ejb-jar-binding>
      <jar-name>yourEjb20.jar</jar-name>
      <connection-factory>
        <jndi-name>eis/jdbc/YourData_CMP</jndi-name>
        <res-auth>Container</res-auth>
      </connection-factory>
    </ejb-jar-binding>
  </module-bindings>
</dfltbndngs>
```

### Setting the connection factory binding for an EJB file

```
<?xml version="1.0">
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
    <ejb-jar-binding>
      <jar-name>yourEjb20.jar</jar-name>
      <ejb-bindings>
        <ejb-binding>
          <ejb-name>YourCmp20</ejb-name>
          <!-- this matches the ejb-name tag in the DD -->
          <connection-factory>
            <jndi-name>eis/jdbc/YourData_CMP</jndi-name>
            <res-auth>PerConnFact</res-auth>
          </connection-factory>
        </ejb-binding>
      </ejb-bindings>
    </ejb-jar-binding>
  </module-bindings>
</dfltbndngs>
```

### Overriding a Resource Ref Binding from a WAR, EJB JAR file, or J2EE client JAR file

Example code for overriding a Resource Ref Binding from a WAR file follows. Use similar code to override a Resource Ref Binding from an EJB JAR file or a J2EE client JAR file.

```
<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
    <war-binding>
      <jar-name>hello.war</jar-name>
      <resource-ref-bindings>
        <resource-ref-binding>
          <!-- the following must match the resource-ref in the DD -->
          <resource-ref-name>jdbc/MyDataSrc</resource-ref-name>
          <jndi-name>war/override/dataSource</jndi-name>
        </resource-ref-binding>
      </resource-ref-bindings>
    </war-binding>
  </module-bindings>
</dfltbndngs>
```

### Overriding MDB JMS listener ports

```
<?xml version="1.0"?>
<!DOCTYPE dfltbndngs SYSTEM "dfltbndngs.dtd">
<dfltbndngs>
  <module-bindings>
```

```

    <ejb-jar-binding>
      <jar-name>YourEjbJar.jar</jar-name>
      <ejb-bindings>
        <ejb-binding>
          <ejb-name>YourMDB</ejb-name>
          <listener-port>yourMdbListPort</listener-port>
        </ejb-binding>
      </ejb-bindings>
    </ejb-jar-binding>
  </module-bindings>
</df1tbdngs>

```

## Example: Installing an EAR file using the default bindings

An example of a simple .ear file installation using the default bindings follows:

1. Go to the Preparing for application install pages. Click **Applications > Install an Application** in the console navigation tree.
2. For **Path**, specify the full path name of the .ear file. For this example, the base file name is my\_appl.ear and the file resides on a server at C:\sample\_apps.
  - a. Select the **Server path** radio button and click **Browse**.
  - b. On the Browse Remote Filesystems page, click on the name of the node that holds the my\_appl.ear file, C:\, **sample\_apps**, **my\_appl.ear**, and then **OK**.
3. Now that a value is given for **Path**, on the first Preparing for application install page, click **Next**.
4. On the second Preparing for application install page, place a checkmark beside the **Generate Default Bindings** check box and click **Next**. Using the default bindings causes any incomplete bindings in the application to be filled in with default values. Existing bindings are not changed. By choosing this option, you can directly jump to the Summary step.
5. On the Install New Application page, click on **Summary**, the last step.
6. On the Summary panel, verify the cell, node, and server onto which the application files will install.
  - a. Beside the **Cell/Node/Server** option, click **Click here**.
  - b. On the **Map modules to application servers** panel, select the server onto which the application files will install from the **Clusters and Servers** list, place a checkmark in the check box beside **Module** to select all of the application modules, and click **Next**.

Because my\_appl.ear does not require any additional settings to complete an installation, the Summary panel displays again.

7. On the Summary panel, click **Finish**.

---

## Enterprise application collection

Use this page to view and manage enterprise applications.

To view this administrative console page, click **Applications > Enterprise Applications**.

### Name

Specifies the name of the installed (or deployed) application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names.

## Status

Indicates whether the application deployed on the application server is started, stopped, or unavailable.

## Enterprise application settings

Use this page to configure an enterprise application.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name***.

### Name

Specifies a logical name for the application. Application names must be unique within a cell and cannot contain characters that are not allowed in object names.

Data type                      String

### Starting Weight

Specifies the order in which applications are started when the server starts. The application with the lowest starting weight is started first.

Data type                      Integer  
Default                          1  
Range                            0 to 100

### Application Binaries

Specifies the directory to which the application EAR file will be installed. The default value is the value of `APP_INSTALL_ROOT/cell_name`, where the `APP_INSTALL_ROOT` variable is `install_root/installedApps`; for example, `C:\WebSphere\AppServer\installedApps\cell_name`.

You can specify an absolute path or use a pathmap variable such as `${MY_APPS}`. You can use a pathmap variable in any installation though it is particularly needed when installing an application on a cluster with members on heterogeneous nodes because, in such cases, there might not be a single way to specify an absolute path. A WebSphere Application Server variable `{CELL}` that denotes the current cell name can also be in the pathmap variable; for example, `${MY_APP}/${CELL}`.

Data type                      String  
Units                            Full path name

### Use Metadata From Binaries

Specifies whether the application server uses the binding, extensions, and deployment descriptors located with the application deployment document, the `deployment.xml` file (default), or those located in the enterprise application resource (EAR) file.

Data type                      Boolean  
Default                          true

### Enable Distribution

Specifies whether WebSphere Application Server expands or deletes application binaries in the installation destination. The default is to enable application distribution. Application binaries for installed applications are expanded to the directory specified. The binaries are also deleted when you uninstall and save

changes to the configuration. If you disable this option, then you must ensure that the application binaries are expanded appropriately in the destination directories of all nodes where the application runs.

|           |         |
|-----------|---------|
| Data type | Boolean |
| Default   | true    |

### **ClassLoader Mode**

Specifies whether the class loader searches in the parent class loader or in the application class loader first to load a class. The standard for JDK class loaders and WebSphere Application Server class loaders is PARENT\_FIRST. By specifying PARENT\_LAST, your application can override classes contained in the parent class loader, but this action can potentially result in ClassCastException or LinkageErrors if you have mixed use of overridden classes and non-overridden classes.

The options are PARENT\_FIRST and PARENT\_LAST. The default is to search in the parent class loader before searching in the application class loader to load a class.

|           |              |
|-----------|--------------|
| Data type | String       |
| Default   | PARENT_FIRST |

### **WAR Classloader Policy**

Specifies whether to use a single class loader to load all WAR files of this application or to use a different class loader for each WAR file.

The options are APPLICATION and MODULE. The default is to use a separate class loader to load each WAR file.

|           |        |
|-----------|--------|
| Data type | String |
| Default   | MODULE |

### **Create MBeans for Resources**

Specifies whether to create MBean files for various resources (such as servlets or JSP files) within an application.

|           |         |
|-----------|---------|
| Data type | Boolean |
| Default   | true    |

### **Reload Enabled**

Specifies whether to enable class reloading when application files are updated.

|           |         |
|-----------|---------|
| Data type | Boolean |
| Default   | true    |

### **Reload Interval**

Specifies the number of seconds to scan the application's file system for updated files. The default is the value of the reload interval attribute in the IBM extension (META-INF/ibm-application-ext.xml) file of the EAR file. This setting takes effect only if class reloading is enabled.

The reload interval specified here overrides the value specified in the IBM extensions for each Web module in the EAR file (which in turn overrides the reload interval specified in the IBM extensions for the application in the EAR file).



|           |         |
|-----------|---------|
| Data type | Integer |
| Units     | Seconds |

### Target mapping collection

Use this page to manage mappings of deployed applications or modules to servers or clusters.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name* > Target Mappings**.

**Target:** States the name of the target server or cluster to which the application or module maps. You specify the target on the Map modules to application servers page accessed from the settings for an application.

**Node:** Specifies the node name if the target is a server.

**Status:** Indicates whether the status of the application running on the target server or cluster is started, stopped or unavailable.

### Target mapping settings

Use this page to map a deployed application or module to a server or cluster.

To view this administrative console page, click **Applications > Enterprise Applications > *application\_name* > Target Mappings > *target\_name***.

**Target:** States the name of the target server or cluster to which the application or module maps. You specify the target on the Map modules to application servers page accessed from the settings for an application.

|           |        |
|-----------|--------|
| Data type | String |
|-----------|--------|

**Enabled:** Indicates whether the application modules installed on the target server are started (or enabled) when the server starts. This sets the initial state of application modules. A true value indicates that the corresponding modules are enabled and thus are accessible when the server starts. A false value indicates that the corresponding modules are not enabled and thus are not accessible when the server starts.

|           |         |
|-----------|---------|
| Data type | Boolean |
| Default   | true    |

---

## Starting and stopping applications

You can start an application that is not running (has a status of *Stopped*) or stop an application that is running (has a status of *Started*).

### Steps for this task

1. Go to the Enterprise Applications page. Click **Applications > Enterprise Applications** in the console navigation tree.
2. Check the check box for the application you want started or stopped.
3. Click a button:

**Start** Runs the application and changes the state of the application from *Stopped* to *Started*.

**Stop** Stops the processing of the application and changes the state of the application from *Started* to *Stopped*.

To restart a running application, place a checkmark in the check box for the application you want to restart, click **Stop** and then click **Start**.

### Results

The status of the application changes and a message stating that the application started or stopped displays at the top the page.

---

## Exporting applications

You can export an enterprise application to a location of your choice. Exporting applications enables you to back up your applications and preserve binding information for the applications. You might export your applications before updating installed applications or migrating to a later version of the WebSphere Application Server product.

### Steps for this task

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. Place a checkmark in the check box beside the application and click **Export**.
3. On the Export Application EAR Files page, click on the link to download the exported EAR file.
4. Use the browser dialogue to specify a location at which to save the exported EAR file and click **OK**.

The file containing binding information is exported to the specified node and directory, and has the name *enterprise\_application\_name.ear*.

---

## Exporting DDL files

You can export the DDL files (Table.ddl) in the EJB modules of the application to a location of your choice.

### Steps for this task

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. Place a checkmark in the check box beside the application and click **Export DDL**. If the application has no DDL files in any of its EJB modules, then the message *No DDL files were found* displays at the top of the page. If the application has DDL files in its EJB modules, then a page displays listing DDL files in the format *appname.ear/\_module.jar\_Table.ddl*.
3. Click on a file in the list to download the file to your machine.

---

## Updating applications

You can update an application deployed on a server. The steps below describe how to update a deployed application using the administrative console.

### Steps for this task

1. Update the contents of the application and reassemble it, using the Application Assembly Tool. Typical tasks include adding or editing assembly properties, adding or importing modules into an application, and adding enterprise beans, Web components, and files.
2. Go to the Applications page of the administrative console. Click **Applications > Enterprise Applications** in the console navigation tree.
3. **(Optional)** Back up the application. Place a checkmark in the check box beside the application you want uninstalled and click **Export** to export the application to an EAR file and preserve the binding information.
4. With a checkmark beside the application, click **Update**.  
The binding information of the updated (new) version of the application merges with the binding information from the installed (old) version. Then, the older version uninstalls from the configuration and the new version installs.
5. Complete the steps in the Preparing for application install page and the pages that follow it. See information on installing applications and on the settings page for application installation for guidance.  
Note that the installation steps have the merged binding information from the new version and the old version. If the new version has bindings for application artifacts such as EJB JNDI names, EJB references or resource references, then those bindings will be part of the merged binding information. If new bindings are not present, then bindings are taken from the installed (old) version. If bindings are not present in the old version and if the default binding generation option is enabled, then the default bindings will be part of the merged binding information.  
You can select whether to ignore bindings in the old version or ones in the new version.
6. **(Optional)** Map the installed application or module to servers or clusters. Use the Map modules to application servers page of the Install New Application pages displayed during updating the application. Or, after updating the application, use the Map modules to application servers page accessed from the Enterprise Applications page.
  - a. Go to the Map modules to application servers page. Click **Applications > Enterprise Applications** in the console navigation tree, click the application name, and then click **Map modules to application servers**.
  - b. Specify the application server where you want to install modules contained in your application and click **OK**.
7. Click **Save** on the administrative console taskbar to save the changes to your configuration. In the single server (base) product, after you click **Save** the old version of the application is uninstalled and the new version is installed into the configuration. The application binaries for the old version are deleted from the destination directory and the new binaries are copied to the directory. In the Network Deployment product, the old application files are deleted and new files are copied when the configuration on the deployment manager synchronizes with the configuration on the node where the application is installed.  
If the application is running when you update it, the application stops running before its files are copied to the destination directory of the node and restarts after the copy operation completes. Thus, the application is unavailable on the node during the time the node is synchronizing its configuration with the deployment manager.
8. Restart the application so the changes take effect. If the application is updated while it is running, WebSphere Application Server stops the application, updates the application logic and restarts the application.

- a. Click **Applications > Enterprise Applications** in the console navigation tree to go to the Enterprise Applications page.
- b. Check the check box for the updated application.
- c. Click **Start**.

Note that you can also update applications using the wsadmin tool, which provides updating capabilities identical to those provided by clicking **Update** on the Enterprise Applications page. Further, in some situations, you can update applications without needing to restart the applications.

## Hot deployment and dynamic reloading

You can make various changes to applications and their contents without having to stop the server and start it again. Making these types of changes is known as *hot deployment and dynamic reloading*.

Hot deployment is the process of adding new components (such as WAR files, EJB Jar files, enterprise Java beans, servlets, and JSP files) to a running server without having to stop the application server process and start it again.

Dynamic reloading is the ability to change an existing component without needing to restart the server in order for the change to take effect. Dynamic reloading involves:

- Changes to the implementation of a component of an application, such as changing the implementation of a servlet
- Changes to the settings of the application, such as changing the deployment descriptor for a Web module

### Steps for this task

1. Locate your expanded application files. The application files are in the directory you specified when installing the application or, if you did not specify a custom target directory, are in the default target directory, *install\_root/installedApps/cell\_name*. Your EAR file, `${APP_INSTALL_ROOT}/cell_name/application_name.ear`, points to the target directory. The variables.xml file for the node defines `${APP_INSTALL_ROOT}`.

It is important to locate the expanded application files because, as part of installing applications, a WebSphere application server unjars portions of the EAR file onto the file system of the computer that will run the application. These expanded files are what the server looks at when running your application.

If you cannot locate the expanded application files, look at the `binariesURL` attribute in the deployment.xml file for your application. The attribute designates the location the run time uses to find the application files.

For the remainder of this information on hot deployment and dynamic reloading, *application\_root* represents the root directory of the expanded application files.

2. Locate application metadata files. The metadata files include the deployment descriptors (`web.xml`, `application.xml`, `ejb-jar.xml`, and the like), the bindings files (`ibm-web-bnd.xmi`, `ibm-app-bnd.xmi`, and the like), and the extensions files (`ibm-web-ext.xmi`, `ibm-app-ext.xmi`, and the like).

Metadata XML files for an application can be loaded from one of two locations. The metadata files can be loaded from the same location as the application binary files (such as *application\_root/META-INF*) or they can be loaded from the WebSphere configuration tree, `${CONFIG_ROOT}/cells/cell_name/applications/`

*application\_EAR\_name/deployments/application\_name/*. The value of the `useMetadataFromBinary` flag specified during application installation controls which location is used. If specified, the metadata files are loaded from the same location as the application binary files. If not specified, the metadata files are loaded from the application deployment folder in the configuration tree.

For the remainder of this information, *metadata\_root* represents the location of the metadata files for the specified application or module.

3. CAUTION: If you are running WebSphere Application Server on a group of machines using Network Deployment and you are changing an application on a particular node, disable automatic synchronization.
  - a. Click **System Administration > Node Agents** in the administrative console navigation tree, click on a node agent name, and then click **File Synchronization Service**.
  - b. On the File Synchronization Service page, remove the checkmark from the check box for **Automatic Synchronization** and click **OK**.

When you run WebSphere Application Server on a group of machines using Network Deployment and you change a file on the disk in the expanded application directory for a particular node, you can lose those changes the next time node synchronization occurs. In the Network Deployment environment, the configuration stored by the deployment manager is the master copy and any changes detected between that master copy and the copy on a particular machine trigger the master copy to be downloaded to the node.

4. Change or add the following components or modules as needed:
  - Application files
  - WAR files
  - EJB Jar files
  - HTTP plug-in configuration files
5. For changes to take effect, you might need to start, stop, or restart an application. "Starting and stopping applications" provides information on using the administrative console to start, stop, or restart an application. "Example: Starting an application using wsadmin" (not in this document) and "Example: Stopping running applications on a server using wsadmin" (not in this document) provide information on using the wsadmin scripting tool.
6. If you disabled automatic synchronization in step 3, return to the File Synchronization Service page, enable **Automatic Synchronization**, and click **OK**.

## Changing or adding application files

You can change or add application files on application servers without having to stop the server and start it again. This file describes—

- Updating an existing application on a running server (providing a new EAR file)
- Adding a new application to a running server
- Removing an existing application from a running server
- Adding a new EJB or Web module to an existing, running application
- Changing the application.xml file for an application
- Changing the ibm-app-ext.xmi file for an application
- Changing the ibm-app-bnd.xmi file for an application
- Changing a non-module Jar file contained in the EAR file

## Updating an existing application on a running server (providing a new EAR file)

Reinstall an updated application using the administrative console or the wsadmin \$AdminApp install command with the -update option.

Both reinstallation methods enable you to update an existing application using any of the other steps listed in this file, including changing classes, adding modules, removing modules, changing modules, or changing metadata files. The application reinstallation methods detect the changes in your application and prompt you for additional binding data that might be needed to install the application. The reinstallation process automatically stops and restarts your application on the appropriate servers.

Hot deployment:           Yes  
Dynamic reloading:       Yes

### **Adding a new application to a running server**

Install an application using the administrative console or the wsadmin install command.

Hot deployment:           Yes  
Dynamic reloading:       No

### **Removing an existing application from a running server**

Stop the application and then uninstall it from the server. Use the administrative console to stop the application and then uninstall it. Or run the wasadmin stopApplication command and then the uninstall command.

Hot deployment:           Yes  
Dynamic reloading:       No

### **Adding a new EJB or Web module to an existing, running application**

1. Update the application files in the *application\_root* location.
2. Restart the application. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:           Yes  
Dynamic reloading:       No

### **Changing the application.xml file for an application**

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment:           Not applicable  
Dynamic reloading:       Yes

### **Changing the ibm-app-ext.xmi file for an application**

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the wasadmin stopApplication and startApplication commands.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Changing the `ibm-app-bnd.xmi` file for an application

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Changing a non-module Jar file contained in the EAR file

1. Update the non-module Jar file in the `application_root` location.
2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.  
If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Hot deployment: Yes  
Dynamic reloading: Yes

### Changing or adding WAR files

You can change WAR files on application servers without having to stop the server and start it again. This file describes—

- Changing an existing JSP file
- Adding a new JSP file to an existing application
- Changing an existing servlet class (editing and recompiling)
- Changing a dependent class of an existing servlet class
- Adding a new servlet using the Invoker (Serve Servlets by class name) facility or adding a dependent class to an existing application
- Adding a new servlet, including a new definition of the servlet in the `web.xml` deployment descriptor for the application
- Changing the `web.xml` file of a WAR file
- Changing the `ibm-web-ext.xmi` file of a WAR file
- Changing the `ibm-web-bnd.xmi` file of a WAR file

### Changing an existing JSP file

Place the changed JSP file directly in the `application_root/module_name` directory or the appropriate subdirectory. The change will be automatically detected and the JSP will be recompiled and reloaded.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Adding a new JSP file to an existing application

Place the new JSP file directly in the *application\_root/module\_name* directory or the appropriate subdirectory. The new file will be automatically detected and compiled on the first request to the page.

Hot deployment: Yes  
Dynamic reloading: Yes

### Changing an existing servlet class (editing and recompiling)

1. Place the new version of the servlet .class file directly in the *application\_root/module\_name/WEB-INF/classes* directory. If the .class file is part of a Jar file, you can place the new version of the Jar file directly in *application\_root/module\_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.
2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Changing a dependent class of an existing servlet class

1. Place the new version of the dependent .class file directly in the *application\_root/module\_name/WEB-INF/classes* directory. If the .class file is part of a Jar file, you can place the new version of the Jar file directly in *application\_root/module\_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.
2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Adding a new servlet using the Invoker (Serve Servlets by class name) facility or adding a dependent class to an existing application

1. Place the new .class file directly in the *application\_root/module\_name/WEB-INF/classes* directory. If the .class file is part of a Jar file, you can place the new version of the Jar file directly in *application\_root/module\_name/WEB-INF/lib*. In either case, the change will be detected, the Web application will be shut down and reinitialized, picking up the new class.

This case is treated the same as changing an existing class. The difference is that adding the servlet or class does not immediately cause the Web application to reload because the class has never been loaded before. The class simply becomes available for execution.



2. If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.

Hot deployment: Yes  
Dynamic reloading: Not applicable

### **Adding a new servlet, including a new definition of the servlet in the web.xml deployment descriptor for the application**

1. Place the new .class file directly in the *application\_root/module\_name/WEB-INF/classes* directory. If the .class file is part of a Jar file, you can place the new version of the Jar file directly in *application\_root/module\_name/WEB-INF/lib*.

You can edit the web.xml file in place or copy it into the *application\_root/module\_name/WEB-INF/classes* directory. The new .class file will not trigger a reloading of the application.

2. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands. After the application restarts, the new servlet is available for service.

Hot deployment: Yes  
Dynamic reloading: Not applicable

### **Changing the web.xml file of a WAR file**

1. Edit the web.xml file in place or copy it into the *metadata\_root/module\_name/WEB-INF* directory.
2. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Yes  
Dynamic reloading: Yes

### **Changing the ibm-web-ext.xmi file of a WAR file**

Edit the extension settings as needed. You can change all of the extension settings. The only warning is if you set the `reloadInterval` property to zero (0) or the `reloadEnabled` property to false, the application will no longer automatically detect changes to class files. Both of these changes disable the automatic reloading function. The only way to re-enable automatic reloading is to change the appropriate property and restart the application. See other task descriptions in this file for information on restarting an application.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### **Changing the ibm-web-bnd.xmi file of a WAR file**

1. Edit the bindings as needed. You can change all of the values but ensure that the entities you are binding to are present in the configuration of the server.
2. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Changing or adding EJB Jar files

You can change EJB Jar files on application servers without having to stop the server and start it again. This file describes—

- Changing the `ejb-jar.xml` file of an EJB Jar file
- Changing the `ibm-ejb-jar-ext.xmi` or `ibm-ejb-jar-bnd.xmi` file of an EJB Jar file
- Changing the `Table.ddl` file for an EJB Jar file
- Changing the `Map.mapxmi` or `Schema.dbxmi` file for an EJB Jar file
- Updating the implementation class for an EJB file or a dependent class of the implementation class for an EJB file
- Updating the Home/Remote interface class for an EJB file
- Adding a new EJB file to an existing EJB Jar file

### Changing the `ejb-jar.xml` file of an EJB Jar file

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Changing the `ibm-ejb-jar-ext.xmi` or `ibm-ejb-jar-bnd.xmi` file of an EJB Jar file

Restart the application. Automatic reloading will not detect the change. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Changing the `Table.ddl` file for an EJB Jar file

Rerun the DDL file on the user database server. Changing the `Table.ddl` file has no effect on the application server and is a change to the database table schema for the EJB files.

Hot deployment: Not applicable  
Dynamic reloading: Not applicable

### Changing the `Map.mapxmi` or `Schema.dbxmi` file for an EJB Jar file

1. Change the `Map.mapxmi` or `Schema.dbxmi` file for an EJB Jar file.
2. Regenerate the deployed code artifacts for the EJB file.
3. Apply the new EJB Jar file to the server.
4. Restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Updating the implementation class for an EJB file or a dependent class of the implementation class for an EJB file

1. Update the class file in the *application\_root/module\_name.jar* file.
2. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.  
If automatic reloading is not enabled, restart the application of which the EJB file is a member. If the updated module is used by other modules in other applications, restart those applications as well. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Updating the Home/Remote interface class for an EJB file

1. Update the interface class of the EJB file.
2. Regenerate the deployed code artifacts for the EJB file.
3. Apply the new EJB Jar file to the server.
4. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.  
If automatic reloading is not enabled, restart the application of which the EJB file is a member. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Not applicable  
Dynamic reloading: Yes

### Adding a new EJB file to an existing EJB Jar file

1. Apply the new or updated Jar file to the *application\_root* location.
2. If automatic reloading is enabled, you do not need to take further action. Automatic reloading will detect the change.  
If automatic reloading is not enabled, restart the application. Use the administrative console to restart the application. Or run the `wasadmin stopApplication` and `startApplication` commands.

Hot deployment: Yes  
Dynamic reloading: Yes

### Changing the HTTP plug-in configuration

You can change the HTTP plug-in configuration without having to stop the server and start it again. This file describes—

- Changing the `application.xml` file to change the context root of a WAR file
- Changing the `web.xml` file to add, remove, or modify a servlet mapping
- Changing the `server.xml` file to add, remove, or modify an HTTP transport or changing the `virtualhost.xml` file to add or remove a virtual host or to add, remove, or modify a virtual host alias

### Changing the `application.xml` file to change the context root of a WAR file

1. Change the `application.xml` file.

2. Regenerate the plug-in configuration file using the administrative console or by running the GenPluginCfg.bat/sh script.

Hot deployment: Yes  
Dynamic reloading: No

### Changing the web.xml file to add, remove, or modify a servlet mapping

1. Change the web.xml file.
2. Regenerate the plug-in configuration file using the administrative console or by running the GenPluginCfg.bat/sh script.

If the Web application has file serving enabled or has a servlet mapping of /, you do not have to regenerate the plug-in configuration. In all other cases the regeneration is required.

Hot deployment: Yes  
Dynamic reloading: Yes

### Changing the server.xml file to add, remove, or modify an HTTP transport or changing the virtualhost.xml file to add or remove a virtual host or to add, remove, or modify a virtual host alias

1. Change the server.xml file to add, remove, or modify an HTTP transport or change the virtualhost.xml file to add or remove a virtual host or to add, remove, or modify a virtual host alias.
2. Regenerate the plug-in configuration file using the administrative console, by running the GenPluginCfg.bat/sh script, or by running a wsadmin command.

Hot deployment: Yes  
Dynamic reloading: Yes

---

## Uninstalling applications

After an application no longer is needed, you can uninstall it. Uninstalling an application deletes the application from the WebSphere Application Server configuration repository and it deletes the application binaries from the file system of all nodes where the application modules are installed.

### Steps for this task

1. Click **Applications > Enterprise Applications** in the administrative console navigation tree to access the Enterprise Applications page.
2. **(Optional)** Stop the application. Place a checkmark in the check box beside the application you want uninstalled and click **Stop**. Note that it is recommended you stop the application before uninstalling. However, even if the application is running when uninstallation commences, the uninstallation program stops the application before deleting application binaries from the file system.
3. **(Optional)** Back up the application. Place a checkmark in the check box beside the application you want uninstalled and click **Export** to export the application to an EAR file and preserve the binding information.
4. With a checkmark in the check box beside the application you want uninstalled, click **Uninstall**.
5. Confirm the uninstallation operation.

6. Click **Save** on the console taskbar to save changes made to the administrative configuration.

In the single-server (base) product, application binaries are deleted after you click **Save**. In the Network Deployment product, application binaries are deleted when configuration changes on the deployment manager synchronize with configurations for individual nodes.

---

## Deploying and managing applications: Resources for learning

Use the following links to find relevant supplemental information about deploying and managing applications using the administrative console. The information resides on IBM and non-IBM Internet sites, whose sponsors control the technical accuracy of the information.

These links are provided for convenience. Often, the information is not specific to the IBM WebSphere Application Server product, but is useful all or in part for understanding the product. When possible, links are provided to technical papers and Redbooks that supplement the broad coverage of the release documentation with in-depth examinations of particular product areas.

View links to additional information about:

- Programming model and decisions
- Programming instructions and examples
- Administration

### Programming model and decisions

-  **The J2EETM Tutorial: The Duke's Bank Application**  
([http://java.sun.com/j2ee/tutorial/1\\_3-fcs/doc/Ebank.html](http://java.sun.com/j2ee/tutorial/1_3-fcs/doc/Ebank.html))
-  **Best Practices in WebSphere Application: Separating the developers from the administrators** (<http://www.sys-con.com/websphere/articleprint.cfm?id=26>)
-  **Designing Enterprise Applications with the Java™ 2 Platform, Enterprise Edition, Second Edition** ([http://java.sun.com/blueprints/guidelines/designing\\_enterprise\\_applications\\_2e/](http://java.sun.com/blueprints/guidelines/designing_enterprise_applications_2e/))
-  **Designing Enterprise Applications, Second Edition**  
(<http://developer.java.sun.com/developer/Books/j2ee/designingenterprise/>)
-  **Building Java™ Enterprise Applications Volume I: Architecture**  
(<http://developer.java.sun.com/developer/Books/j2ee/bjeapps/>)

### Programming instructions and examples

-  **WebSphere Application Server education**  
(<http://www.ibm.com/software/webservers/learn/>)
-  **Developing and Testing a Complete 'Hello World' J2EE Application with IBM WebSphere Studio Application Developer for Linux**  
([http://www7b.software.ibm.com/wsdd/library/tutorials/0206\\_wosnick/wosnick\\_reg.html?open&l=937,t=gr](http://www7b.software.ibm.com/wsdd/library/tutorials/0206_wosnick/wosnick_reg.html?open&l=937,t=gr))
-  **Writing Enterprise Applications with Java™ 2 Platform, Enterprise Edition** (<http://developer.java.sun.com/developer/onlineTraining/J2EE/Intro/>)

### Administration

- ✓ **Listing of all IBM WebSphere Application Server Redbooks**  
(<http://publib-b.boulder.ibm.com/Redbooks.nsf/Portals/WebSphere>)