

IBM WebSphere InterChange Server



# System Administration Guide

*V4.2.2*

**Note!**

Before using this information and the product it supports, read the information in "Notices" on page 185.

**20March2004**

This edition of this document applies to IBM WebSphere InterChange Server, version 4.2.2, and to all subsequent releases and modifications until otherwise indicated in new editions.

To send us your comments about this document, e-mail [doc-comments@us.ibm.com](mailto:doc-comments@us.ibm.com). We look forward to hearing from you.

When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© Copyright International Business Machines Corporation 1997, 2004. All rights reserved.

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

---

# Contents

<b>About this document</b> . . . . .	<b>v</b>
Audience . . . . .	v
Related documents . . . . .	v
Typographic conventions . . . . .	v
<b>New in this release</b> . . . . .	<b>vii</b>
New in release 4.2.2 . . . . .	vii
New in release 4.2.1 . . . . .	viii
New in release 4.2.0 . . . . .	viii
New in release 4.1.1. . . . .	ix
New in release 4.1.0. . . . .	ix
New in release 4.0.1. . . . .	ix
New in release 4.0.0 . . . . .	x
<b>Chapter 1. Monitoring the system.</b> . . . .	<b>1</b>
Using System Monitor . . . . .	1
Using the monitoring functions of InterChange Server Component Management view of System Manager . . . . .	22
Using the SNMP agent . . . . .	28
Using persistent monitoring . . . . .	39
Using WebSphere Business Integration Monitor . . . . .	40
<b>Chapter 2. Operating components of the system.</b> . . . .	<b>47</b>
Overview of operating the system . . . . .	47
Operating InterChange Server . . . . .	48
Using System Manager . . . . .	52
Operating connectors . . . . .	56
Operating collaboration objects . . . . .	69
Operating maps . . . . .	76
Operating relationships . . . . .	78
Using Relationship Manager. . . . .	82
Backing up IBM WebSphere InterChange Server components. . . . .	96
Using Repos_Copy . . . . .	99
Scheduling jobs. . . . .	110
<b>Chapter 3. Troubleshooting the system</b> . . . . .	<b>117</b>
Using log and trace files for troubleshooting . . . . .	117
Controlling server memory usage . . . . .	143
Maintaining event archives . . . . .	145
Managing WIP connections. . . . .	145
Managing failures . . . . .	145
Using Flow Manager . . . . .	152
Resolving collaboration deadlocks . . . . .	162
Managing WebSphere MQ . . . . .	165
Managing database connection pools . . . . .	166
Managing High-Availability (HA) systems. . . . .	167
Managing runtime properties . . . . .	169
Managing the Object Request Broker . . . . .	170
<b>Appendix A. WebSphere MQ reference</b> . . . . .	<b>177</b>
<b>Appendix B. Requirements for restarting IBM WebSphere Business Integration system components</b> . . . . .	<b>179</b>
InterChange Server restart requirements . . . . .	179
Collaboration template restart requirements . . . . .	179

Collaboration object restart requirements . . . . .	180
Connector restart requirements . . . . .	180
Business object restart requirements . . . . .	181
Map restart requirements . . . . .	182
Relationship restart requirements . . . . .	182
<b>Notices . . . . .</b>	<b>185</b>
Programming interface information . . . . .	186
Trademarks and service marks . . . . .	186
<b>Index . . . . .</b>	<b>189</b>

---

## About this document

IBM<sup>(R)</sup> WebSphere<sup>(R)</sup> InterChange Server and its associated toolset are used with IBM WebSphere Business Integration Adapters to provide business process integration and connectivity among leading e-business technologies and enterprise applications.

This document describes how to monitor, operate, and troubleshoot the WebSphere InterChange Server system.

---

## Audience

This document is for system administrators, consultants, and developers who administer the WebSphere InterChange Server system.

---

## Related documents

The complete set of documentation available with this product describes the features and components common to all WebSphere Integration Server installations, and includes reference material on specific components.

This document contains many reference to the System Installation Guide. If you choose to print this document, you may want to print that guide as well.

You can install the documentation or read it directly online at one of the following sites:

- For InterChange Server documentation:  
<http://www.ibm.com/websphere/integration/wicserver/infocenter>
- For collaboration documentation:  
<http://www.ibm.com/websphere/integration/wbicollaborations/infocenter>
- For WebSphere Business Integration Adapters documentation:  
<http://www.ibm.com/websphere/integration/wbiadapters/infocenter>

These sites contain simple directions for downloading, installing, and viewing the documentation.

---

## Typographic conventions

This document uses the following conventions:

<code>courier font</code>	Indicates a literal value, such as a command name, filename, information that you type, or information that the system prints on the screen.
<b>bold</b>	Indicates a new term the first time that it appears.
<i>italic, italic</i>	Indicates a variable name or a cross-reference.
<i>blue outline</i>	A blue outline, which is visible only when you view the manual online, indicates a cross-reference hyperlink. Click inside the outline to jump to the object of the reference.
{ }	In a syntax line, curly braces surround a set of options from which you must choose one and only one.

---

---

[ ]	In a syntax line, square brackets surround an optional parameter.
...	In a syntax line, ellipses indicate a repetition of the previous parameter. For example, <code>option[,...]</code> means that you can enter multiple, comma-separated options.
< >	In a naming convention, angle brackets surround individual elements of a name to distinguish them from each other, as in <code>&lt;server_name&gt;&lt;connector_name&gt;tmp.log</code> .
/, \	In this document, backslashes (\) are used as the convention for directory paths. For UNIX installations, substitute slashes (/) for backslashes. All IBM WebSphere InterChange Server product pathnames are relative to the directory where the IBM WebSphere InterChange Server product is installed on your system. Paragraphs inside a box with a UNIX or Windows label indicate notes listing operating system differences.
<div style="border: 1px solid black; padding: 2px; display: inline-block;">UNIX/Windows</div>	
<code>%text%</code> and <code>\$text</code>	Text within percent (%) signs indicates the value of the Windows text system variable or user variable. The equivalent notation in a UNIX environment is <code>\$text</code> , indicating the value of the <code>text</code> UNIX environment variable.
<code>ProductDir</code>	Represents the directory where the product is installed.

---

---

## New in this release

---

### New in release 4.2.2

#### February 2004

In the 4.2.2.2 release of the IBM WebSphere InterChange Server product, this guide has been updated with the following changes:

- Revised the section “Using WebSphere Business Integration Monitor” on page 40 with additional information about configuring flow monitoring.
- Made minor revisions to the following sections:
  - “Stopping the SNMP agent” on page 33
  - “Setting automatic and remote restart for a connector” on page 66
  - “Managing High-Availability (HA) systems” on page 167
  - “Managing the Object Request Broker” on page 170

#### December 2003

In the 4.2.2 release of the IBM WebSphere InterChange Server product, this guide has been updated with the following changes:

- Replaced VisiBroker ORB with IBM ORB
- Revised the following sections:
  - “Using Flow Manager” on page 152 to reflect minor GUI (Graphical User Interface) changes
  - “Setting automatic and remote restart for a connector” on page 66 to reflect the change from VisiBroker ORB to IBM ORB
- Added the following sections:
  - “Using Relationship Manager” on page 82
  - “Configuring stack trace” on page 132
  - “Using WebSphere Business Integration Monitor” on page 40
- Made the following System Monitor terminology changes:
  - Removed all instances of “Windows-based System Monitor.” The tool formerly referred to as the Windows-based System Monitor no longer exists as a separate tool. The functionality of the Windows-based System Monitor has been moved to the InterChange Server Component Management view of System Manager. See the section “Using the monitoring functions of InterChange Server Component Management view of System Manager” on page 22 for further details.
  - Removed all instances of “Web-based System Monitor” and replaced them with “System Monitor.” The tool formerly referred to as the Web-based System Monitor is now called System Monitor. There is only one tool named System Monitor.

---

## New in release 4.2.1

In the 4.2.1 release of the IBM WebSphere InterChange Server product, this guide has been updated with the following changes:

- Revised the section “Using System Monitor” on page 1
- Added the section “Using persistent monitoring” on page 39
- Added the section “Preventing collaboration deadlocks” on page 164
- Clarified operating system support for Web-based System Monitor (see “Required software for System Monitor” on page 2)
- Added the following properties to the section “Configuring other InterChange Server logging and tracing parameters” on page 132:
  - SLEEP\_TIME
  - MAX\_TRACE\_WRITE\_TRIES
- Added remote start functionality to the OADAutoRestartAgent standard connector property (see Table 12 on page 67)
- Made editorial improvements, including some reorganization of information

---

## New in release 4.2.0

The 4.2.0 release of the IBM WebSphere InterChange Server product includes the following new features and changes:

- “New features”
- “New organization of this guide” on page ix
- “New terminology” on page ix

### New features

The following new features are covered in this guide:

- **Web-based System Monitor:** The Web-based System Monitor is a new, Web-based tool that allows you to monitor the IBM WebSphere ICS system from the Web. It coexists with the Windows-based System Monitor, which is a revised version of the System Monitor provided in previous releases. For a full description of the differences between these two versions of System Monitor, and for instructions on using each tool, see “Using System Monitor” on page 1.
- **System Manager:** System Manager is the revised tool that was known as CrossWorlds System Manager (CSM) in previous releases (see “New terminology” on page ix). Descriptions of how to use System Manager to perform certain tasks exist throughout this guide, particularly in “Using System Manager” on page 52, but for a thorough description of the revised functionality of System Manager, see the *System Implementation Guide*.
- **Flow Manager:** Flow Manager is a new tool that allows you to manage unresolved flows. For instructions on using Flow Manager, see “Using Flow Manager” on page 152.
- **Flow control:** Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. For instructions on using flow control, see the following sections:
  - “Using default monitors” on page 2
  - “Creating new monitors using the Monitor Definition Wizard” on page 9
  - “Configuring system-wide flow control” on page 54
  - “Configuring flow control for connectors” on page 68
  - “Configuring flow control for collaboration objects” on page 74



- `Repos_copy`: Many new arguments have been added to and removed from the `repos_copy` tool. For more information, refer to “Using `Repos_Copy`” on page 99.
- Optimized InterChange Server recovery: Refer to “Strategies for InterChange Server recovery” on page 146.
- Long-lived business processes: Refer to “Reconfiguring the timeout attribute for long-lived business processing” on page 75.
- Dynamic updates: Refer to Appendix B, “Requirements for restarting IBM WebSphere Business Integration system components,” on page 179.

## New organization of this guide

This guide has been reorganized in the following ways:

- Topics have been consolidated into single sections that describes both conceptual and task-oriented information. Prior to this release, topics were spread out among many chapters, forcing you to jump from chapter to chapter to piece together various sections of a topic.
- The chapters have been renamed and reorganized to reflect the needs of the system administrator.

## New terminology

Two terminology changes are apparent in this guide: the removal of the term “CrossWorlds” and the use of “Web-based” or “Windows-based” when describing System Monitor.

The “CrossWorlds” name is no longer used to describe an entire system or to modify the names of components or tools, which are otherwise mostly the same as before. For example, “CrossWorlds System Manager” is now “System Manager,” and “CrossWorlds InterChange Server” is now “WebSphere InterChange Server.”

System Monitor has two versions in this release: a Web-based version and a Windows-based version. The Web-based version is new in this release, while the Windows-based version is a revised version of the System Monitor that was part of previous releases. For a full description of the differences between these two versions of System Monitor, see “Using System Monitor” on page 1

---

## New in release 4.1.1

This product has been internationalized in the 4.1.1 release. For details, see the following:

- “Locale for `repos_copy` files” on page 110

---

## New in release 4.1.0

This section lists the new installation features in IBM CrossWorlds version 4.1.0 and describes changes made to this guide since its last release (4.0.1).

- Expanded coverage of Troubleshooting topics
- Editorial improvements to the guide

---

## New in release 4.0.1

This section lists the new installation features in IBM CrossWorlds version 4.0.1 and describes changes made to this guide since its last release (4.0.0).

- The Weblogic type 4 driver for MS SQL Server has been replaced with an IBM CrossWorlds branded type 4 driver.
- The Oracle Thin driver will be used in place of the Weblogic type 2 driver for Oracle database connectivity.

Both the IBM CrossWorlds branded driver and the Oracle Thin driver are type 4 drivers. The Weblogic drivers are no longer supported in IBM CrossWorlds version 4.0.1.

---

## New in release 4.0.0

This section lists the new features in IBM CrossWorlds version 4.0.0 and describes changes made to this guide since its last release (3.1.2).

- Asynchronous Recovery
- Connector Agent Parallelism
- Service Call Transport Exception

Additional new features covered in other IBM CrossWorlds guides include:

- Deployment
- Database connection pools
- Java System Installer
- Java STA Installer
- Java Messaging Service (JMS)
- Static Relationship Caching

## New distribution of content of this guide

The content of the *System Administration Guide* has been revised to focus more closely on the tasks and information needed for administering an IBM CrossWorlds system that has already been fully developed and tested and is now operating in a production environment.

Configuration tasks that are more appropriate for development of an IBM CrossWorlds system have been moved from the *System Administration Guide* to a new book, the *System Implementation Guide*.

Topical headings for the redistributed tasks still exist in this release of the *System Administration Guide*. Under the topic headings for each of the redistributed tasks, you will find cross-references to the *System Implementation Guide* or to other guides to which the information has been moved.

---

## Chapter 1. Monitoring the system

Monitoring the overall health of the IBM WebSphere InterChange Server (ICS) system includes monitoring all IBM WebSphere InterChange Server Components, such as connectors and collaboration objects, as well as the connection to all integrated applications. Several tools exist for monitoring the system, including System Monitor, the InterChange Server Component Management view of System Manager, SNMP agent, and persistent monitoring. This chapter covers each of these methods for system monitoring and includes the following topics:

“Using System Monitor”

“Using the monitoring functions of InterChange Server Component Management view of System Manager” on page 22

“Using the SNMP agent” on page 28

“Using persistent monitoring” on page 39

“Using WebSphere Business Integration Monitor” on page 40

---

### Using System Monitor

System Monitor is a tool that allows you to monitor the IBM WebSphere InterChange Server system from the Web. It allows you to configure how you view the data and also allows you to view historical data as well as current data. System Monitor also allows you to start, stop, and pause components. For instructions on starting, stopping, and pausing components, see Chapter 2, “Operating components of the system,” on page 47.

This section describes the various components involved in configuring and using System Monitor, and covers the following topics:

“Required software for System Monitor” on page 2

“Different approaches for using System Monitor” on page 2

“Using default monitors” on page 2

“Creating new monitors using the Monitor Definition Wizard” on page 9

“Logging on to System Monitor” on page 13

“Contents of the System Monitor Web page” on page 15

“Using Views” on page 16

“Fine-tuning System Monitor” on page 18

## Required software for System Monitor

The following software is required for running System Monitor:

### Web servers supported:

A Web application server which supports JSP 1.1 or later, and servlets 2.2 or later, such as IBM WebSphere Application Server (4.0.5 in a Windows, AIX, Solaris, or HP-UX environment, or 5.0.2 in a Windows, AIX or Solaris environment), or Tomcat (using IBM JDK 1.3.1)

### Software required on Web server:

- DB2 Client (if using DB2 for repository database)

### Client browsers supported:

- Microsoft Internet Explorer 5.5 SP2 or higher, with Adobe SVG Viewer 3.0 plug-in
- Netscape 4.7x (only), with Adobe SVG Viewer 3.0 plug-in

(on a Windows 2000 or Windows XP operating system only)

**Note:** For installation instructions on IBM WebSphere InterChange Server 4.2, see the *System Installation Guide for UNIX or for Windows*. For installation instructions on WebSphere Application Server 4.0 or higher or Tomcat 4.x or higher, see the documentation that came with that product.

## Different approaches for using System Monitor

You can use System Monitor by taking either of the following approaches:

- Use the default **monitors** and default **views** provided with the product to monitor the system.
- Create new **monitors** and **views** before monitoring the system.

**Monitors** are definitions of the information you want to view. You can either create new monitors using the Monitor Definition Wizard, a tool launched from System Manager, or you can use the default monitors provided with the Monitor Definition Wizard. See Table 1 on page 3 for a description of the default monitors provided with the Monitor Definition Wizard. **Views** are definitions of the monitor or monitors you want displayed when you monitor the system using System Monitor. You can create and configure views from System Monitor, or you can use the default views provided with System Monitor. See Table 3 on page 16 for a description of the default views included with System Manager. Once you have created the monitors and views you need, or decided to use the default monitors and views, you can then monitor the ICS system from the Web.

## Using default monitors

Several default monitors are provided with the IBM WebSphere ICS product. You can either monitor the system using the default monitors, or you can create new monitors. To gain a better understanding of what information is included in each monitor, refer to the following sections:

Table 1 on page 3 contains a description of the default monitors provided with the Monitor Definition Wizard.

Table 2 on page 6 contains a description of display options listed in Table 1 on page 3.

“Sample display options” on page 6 contains samples of the display options listed in Table 1..

Table 1. Default monitors

Default monitor	Definition	Display options	Available operations when viewing monitor from the Web
System Overview	Overview of the current status of all major components of the system: collaborations, connectors, maps, and relationships	Table tree (a table with expandable nodes in the first column that display more rows)	<ul style="list-style-type: none"> <li>• Start, stop, pause, and shut down a collaboration</li> <li>• Start, stop, and pause a connector</li> <li>• Start and stop a map</li> <li>• Start and stop a relationship</li> </ul>
Collaboration Statistics	Current status and statistics of all collaborations in the system: <ul style="list-style-type: none"> <li>• Status</li> <li>• Start time</li> <li>• Total flows</li> <li>• Successful flows</li> <li>• Failed flows</li> <li>• Events in process</li> <li>• Queued events</li> <li>• Max concurrent events</li> </ul>	Table	Start, stop, pause, and shut down
Connector Statistics	Current status and statistics of all connectors: <ul style="list-style-type: none"> <li>• Status</li> <li>• Start time</li> <li>• Total up time</li> <li>• Business objects received</li> <li>• Business objects sent</li> <li>• Agent status</li> </ul>	Table	<ul style="list-style-type: none"> <li>• Start, stop, and pause</li> <li>• Restart and Shut down</li> </ul>
Map Status	Status of all maps	Table	Start and stop
Relationship Status	Status of all relationships	Table	Start and stop
Server Statistics	Current statistics of the server: the number of failed and successful calls, events, and flows	Stacked bar	None

Table 1. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor from the Web
Database Connections	<p>Current status of database connections:</p> <ul style="list-style-type: none"> <li>• Number of free connections</li> <li>• Number of active connections</li> <li>• Maximum number of connections</li> <li>• Peak number of connections</li> </ul>	Table	None
Message Queues	<p>Current status of message queues:</p> <ul style="list-style-type: none"> <li>• Current depth</li> <li>• Maximum depth configured</li> </ul>	Table	None
Business Objects	<p>Current statistics of the business objects for a particular connector: business objects sent and business objects received</p>	Table	None
Connector Subscriptions	<p>Current statistics of the subscriptions for a particular connector:</p> <ul style="list-style-type: none"> <li>• Collaboration object</li> <li>• Initiator</li> </ul>	Table	None
Collaboration Events	<p>Current statistics of collaboration events, which includes the following information:</p> <ul style="list-style-type: none"> <li>• Events in process</li> <li>• Queued events</li> </ul>	Bar	None
Historical Server Statistics	<p>Server statistics for a specific period of time. Statistical information:</p> <ul style="list-style-type: none"> <li>• Successful calls</li> <li>• Failed calls</li> <li>• Total calls</li> <li>• Successful events</li> <li>• Failed events</li> <li>• Total events</li> <li>• Successful flows</li> <li>• Failed flows</li> <li>• Total flows</li> </ul> <p>Time intervals:</p> <ul style="list-style-type: none"> <li>• Start date</li> <li>• End date</li> </ul>	Bar	None

Table 1. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor from the Web
Historical Server Flows	<p>Flow statistics of the server for a specific period of time at certain time intervals.</p> <p>Statistical information:</p> <ul style="list-style-type: none"> <li>• Successful flows</li> <li>• Failed flows</li> <li>• Total flows</li> </ul> <p>Time intervals:</p> <ul style="list-style-type: none"> <li>• 15 min., 30 min., hourly, daily, weekly, or monthly</li> <li>• Start date</li> <li>• End date</li> </ul>	<ul style="list-style-type: none"> <li>• Line</li> <li>• Stacked bar</li> <li>• Bar</li> </ul>	None
Historical Collaboration Flows Stack	<p>Flow statistics of a particular collaboration for a specific period of time at certain time intervals.</p> <p>Statistics information:</p> <ul style="list-style-type: none"> <li>• Successful flows</li> <li>• Failed flows</li> <li>• Total flows</li> </ul> <p>Time intervals:</p> <ul style="list-style-type: none"> <li>• 15 min., 30 min., hourly, 4 hours, 12 hours, daily, weekly, or monthly</li> <li>• Start date</li> <li>• End date</li> </ul>	Stacked bar	None
Historical Collaboration Flows Line	<p>Flow statistics of a particular collaboration for a specific period of time at certain time intervals.</p> <p>Statistics information:</p> <ul style="list-style-type: none"> <li>• Successful flows</li> <li>• Failed flows</li> <li>• Total flows</li> </ul> <p>Time intervals:</p> <ul style="list-style-type: none"> <li>• 15 min., 30 min., hourly, 4 hours, 12 hours, daily, weekly, or monthly</li> <li>• Start date</li> <li>• End date</li> </ul>	Line	None
Event Rate	Current number of processed events per minute	Meter	None

Table 1. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor from the Web
Flow Control	<p>Current state of collaboration objects and connectors under Flow Control:</p> <ul style="list-style-type: none"> <li>• Buffered events</li> <li>• Max event capacity</li> <li>• Blocked status (does not apply to non-blocking collaboration)</li> <li>• Events pending in database (applies only to non-blocking collaborations)</li> <li>• Saturated status</li> </ul>	Table	None
State Change Log	<p>Current persisted state changes on a component for a specified time period. State change information:</p> <ul style="list-style-type: none"> <li>• Time stamp</li> <li>• State</li> </ul> <p>Time intervals:</p> <ul style="list-style-type: none"> <li>• Start date</li> <li>• End date</li> </ul>	Table	None

Table 2 describes the configurable properties of each display option listed in Table 1 on page 3. To see examples of the display options, see “Sample display options.”

Table 2. Configurable properties of display options

Display option	Build-time configurable properties (using Monitor Definition Wizard)	Runtime configurable properties (from System Monitor)
<ul style="list-style-type: none"> <li>• Table</li> <li>• Table tree</li> </ul>	<ul style="list-style-type: none"> <li>• Columns to display</li> <li>• Order of columns</li> <li>• Number of rows to display</li> </ul>	<ul style="list-style-type: none"> <li>• Font and color settings of the labels and data</li> <li>• Number of rows to display</li> </ul>
<ul style="list-style-type: none"> <li>• Stacked bar</li> <li>• Line</li> <li>• Bar</li> </ul>	None	<ul style="list-style-type: none"> <li>• Font and color settings of the labels and data</li> <li>• Show or hide values</li> </ul>
Meter	Meter threshold	Font and color settings of the labels and data

## Sample display options

The following samples represent the display options you can select when creating monitors in the Monitor Definition Wizard.

**Note:** The data in the examples is not indicative of actual data in an ICS system.



## Connector Statistics



Display  items

1-9 of 9 items

Select all displayed items

<input type="checkbox"/>	Connector	Status	Start Time	Total Up Time	Business Objects Received	Business Objects Sent	Agent Status
<input type="checkbox"/>	CaesarConnector	Running			0	0	Stopped
<input type="checkbox"/>	ClarifyConnector	Running			0	0	Stopped
<input type="checkbox"/>	EMailConnector	Running			0	0	Stopped
<input type="checkbox"/>	JDBCConnector	Running			0	0	Stopped
<input type="checkbox"/>	JTextConnector	Running			0	0	Stopped
<input type="checkbox"/>	PortConnector	Running			0	0	Stopped
<input type="checkbox"/>	SAPConnector	Running			0	0	Stopped
<input type="checkbox"/>	SiebelConnector	Running			0	0	Stopped
<input type="checkbox"/>	XMLConnector	Paused			0	0	Stopped

Sat Mar 15 16:16:21 PST 2003

Figure 1. Table display option

## System Overview



Display  items

<<First <Previous 1-20 of 32 items [Next >](#) Last>>

Select all displayed items

<input type="checkbox"/>	Component	Type	Status	Start Time	Total Up Time
<input type="checkbox"/>	▼ DOCUMENTATION02DEV420Server	Server		Sat Mar 15 03:46:59 PST 2003	12 h 32 min 41 sec
<input type="checkbox"/>	▼ CustomerSync_ClarifyToSAP	Collaboration	Paused	2003.03.15 03:48:58 , PST	12 h 30 min 42 sec
<input type="checkbox"/>	ClarifyConnector	Connector	Running		
<input type="checkbox"/>	PortConnector	Connector	Running		
<input type="checkbox"/>	PortConnector	Connector	Running		
<input type="checkbox"/>	SAPConnector	Connector	Running		
<input type="checkbox"/>	▶ OrderStatusQuery_WebToCaesar	Collaboration	Running	2003.03.15 03:49:05 , PST	12 h 30 min 35 sec
<input type="checkbox"/>	▶ PassThroughTest_PortToCaesar	Collaboration	Running	2003.03.15 03:49:05 , PST	12 h 30 min 35 sec

Figure 2. Table tree display option




Figure 3. Line display option




Figure 4. Bar display option




Figure 5. Stacked bar display option

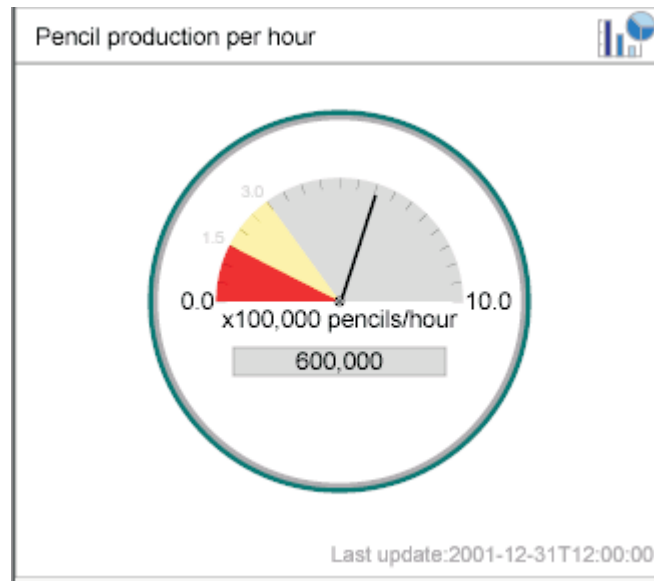


Figure 6. Meter display option

## Creating new monitors using the Monitor Definition Wizard

A monitor is a definition of the information you want to view when monitoring the system using System Monitor. For example, you might have a monitor called System Overview, which displays status and start time of all system components. You create this monitor using the Monitor Definition Wizard, a tool launched from the InterChange Server Component Management view of System Manager.

When you create a monitor, it contains a particular type of system information that can be monitored. Each type of system information has one or more display

options available. Each display option has configurable properties. To see what information types and display options are available when creating a new monitor, see Figure 7.

**Note:** Before creating new monitors, be sure to look at the existing default monitors in Table 1 on page 3, to see if the monitor you want to create already exists.

To create a new monitor, follow these steps:

1. From the InterChange Server Component Management view of System Manager, right-click the server instance you want to connect to, then click Connect. The Server User ID and Password dialog box appears.
2. Type the User ID and password for that server, then click OK. The status of the server changes from "unknown" or "disconnected" to "connected."

**Note:** If the status does not change to "connected," make sure the selected ICS server is running.

3. Right-click the server instance, then select Monitor Definition Wizard. The Monitor Definition Wizard appears. See Figure 7..

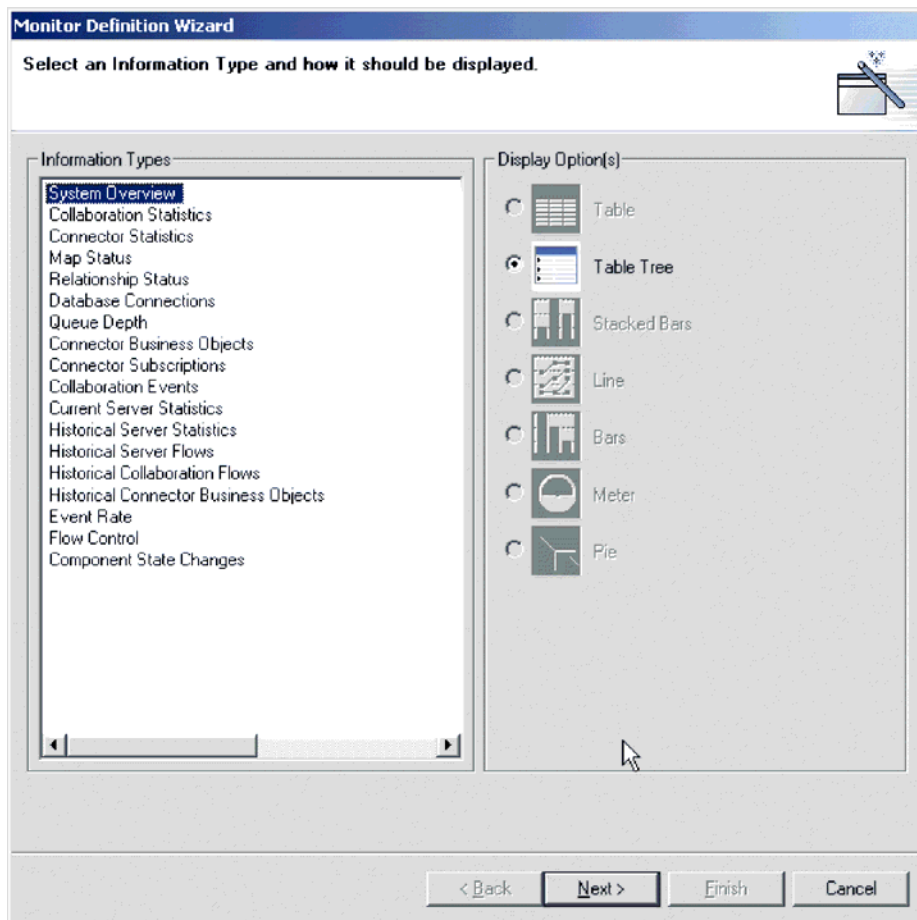


Figure 7. Monitor Definition Wizard, screen for selecting information type and display option

4. Select the type of information you want in the monitor from the Information Types group box, then select how you want the information displayed from the Displayed Option(s) group box. Each information type has one or more

available display options. To see examples of the display options, see “Sample display options” on page 6, and for a description of the configurable properties of each display option, see Table 2 on page 6..

**Note:** If business object probes exist, they will appear in the Information Types group box. For instructions on adding business object probes, refer to the *Collaboration Development Guide*.

**Note:** When you select an information type from the Information Types group box, only the available display options are enabled in the Display Option(s) group box. The unavailable display options remain grayed out.

5. Click Next. The Specify Monitor Properties screen appears (see Figure 8).

Monitor Definition Wizard

**Specify Monitor Properties**

Specify the properties of the Connector Statistics Table monitor.

How should the monitor be described?

Title:  Existing monitors...

Description: (optional)

How should the table look?

Table rows

Number of rows:

Table columns

Choice of Attributes

Select all

Agent Status  
Business Objects Received  
Business Objects Sent  
Start Time  
Status  
Total Up Time

Ordered Attributes

Select all

< Back Next > Finish Cancel

Figure 8. Monitor Definition Wizard, Specify Monitor Properties screen

6. Type a name for the new monitor in the Title field. You can also optionally type a description in the Description field. To make sure you do not use an existing monitor name, click Existing Monitors. This displays a list of the existing monitors.
7. Depending on what information type and display option you chose in the previous screen, you may or may not have more choices to make in this screen. For example, in Figure 8, you can enter the number of rows to appear,

select which attributes to include, and place the chosen attributes in a particular order. These options are available for both Table and Table Tree display options.

Meter is the other display option that has configurable properties in this screen. If you create a monitor with a meter, the configurable property is "meter threshold."

For a list of all display option configurable properties, when using either the Monitor Definition Wizard or System Monitor, see Table 2 on page 6..

8. Depending on which attributes you chose in this screen, either the Next button or the Finish button is enabled.
  - If the Finish button is enabled, then the attributes you chose cannot contain thresholds.
  - If the Next button is enabled, then the attributes you chose can contain thresholds.
9. Click Next or Finish, depending on which button is enabled.
  - When you click Finish, the following message appears, "The monitor was created successfully. Do you want to create another monitor?" Click Yes or No.
  - When you click Next, the Specify Attribute Thresholds screen appears. For an example of a Specify Attribute Thresholds screen, see Figure 9 on page 24. In the Specify Attribute Thresholds screen, you can optionally type a numeric value in the threshold field for each attribute.

**Note:** During runtime, if the value of an attribute exceeds the value of the threshold set for that attribute, the cell that contains the attribute value appears highlighted when it is displayed in the table.

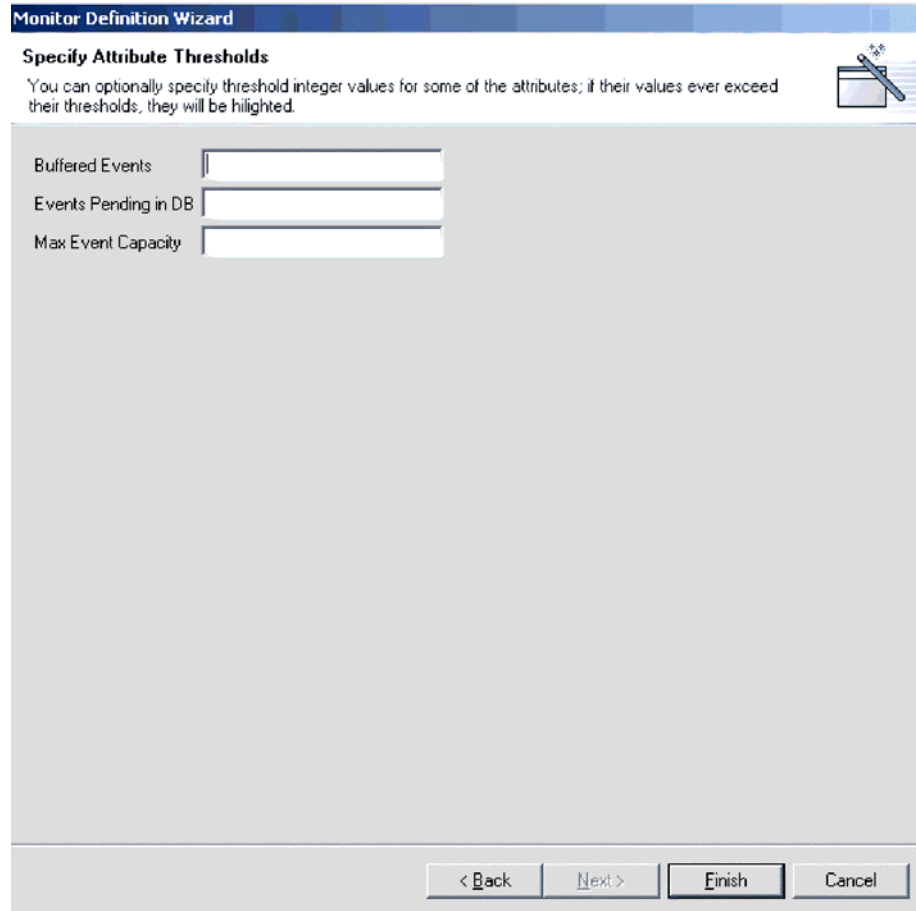


Figure 9. Monitor Definition Wizard, Specify Attribute Thresholds screen

10. After specifying any attribute thresholds, click Finish. The following message appears, "The monitor was created successfully. Do you want to create another monitor?" Click Yes or No.

## Logging on to System Monitor

Once you have either created new monitors or decided to use the default monitors, you are ready to log on to System Monitor to monitor the system. Before you can log on to System Monitor, you must perform the following tasks:

- Start InterChange Server (ICS) on the machine being monitored.
- Make sure System Monitor and the application server are installed. For installation instructions, see the *System Installation Guide for UNIX* or for *Windows*.
- Start the application server.
- Obtain the username and password necessary for logging on to System Monitor. The username and password are the same as those used when logging on to ICS.

To log on to System Monitor, do the following:

1. In a Web browser, navigate to the URL that was specified in the `DASHBOARD_URL` environment variable during installation of System Monitor.  
This URL is made up of three elements:
  - The URL prefix `http://`

- The base URL of your Web server
- The value specified for the context root of the installed application, as specified during the installation of System Monitor. This value must be ICSMonitor.

For instance, if the Web server is named `monitorserver`, and the root context is specified as `ICSMonitor`, as specified in the installation instructions, the URL is: `http://monitorserver/ICSMonitor`

The WebSphere InterChange Server System Monitor login screen appears (see Figure 10).

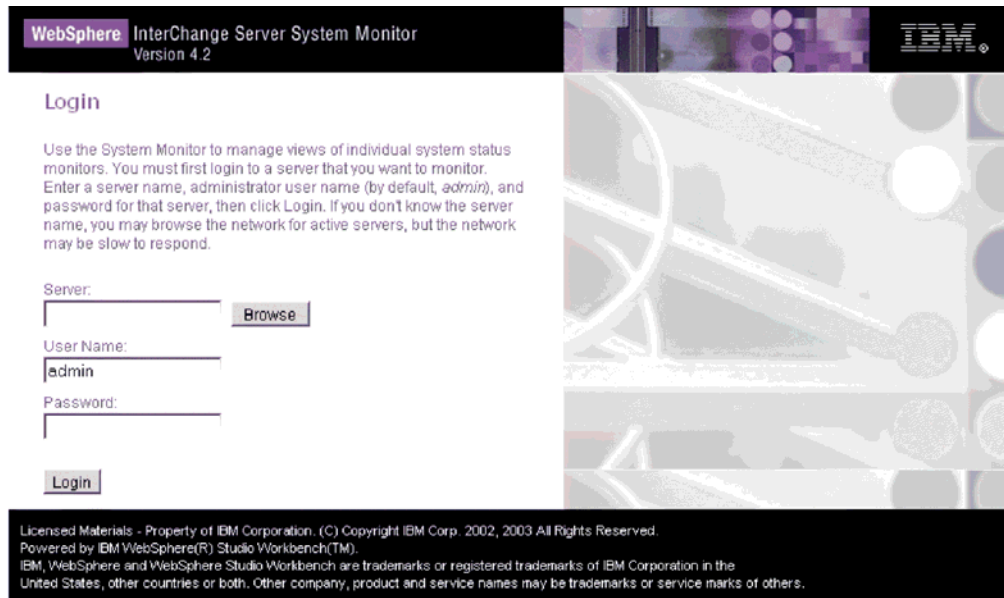


Figure 10. System Monitor, Login screen

2. Type the server name, user name, and password for the InterChange Server instance you want to monitor, then click Login. System Monitor appears (see Figure 11 on page 15).



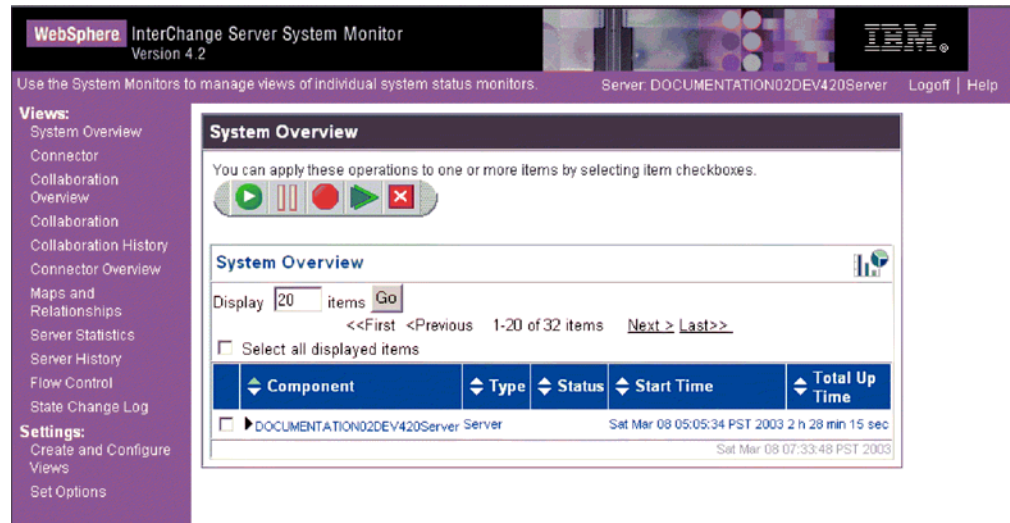


Figure 11. System Monitor, displaying System Overview as the default view

## Contents of the System Monitor Web page

The System Monitor Web page contains the following items:

- **List of views:** Initially, the views listed in the left column under Views are the default views provided with the installation of System Monitor, but you can add, change or delete views to suit your monitoring needs.
- **Create and Configure Views link:** This link opens the Create and Configure Views dialog box (see Figure 12 on page 17), which allows you to create, configure, or delete views. It also allows you to set the default view you see when you log on to System Monitor. For instructions on creating, configuring, or deleting views, or setting the default view you see when you log on, refer to the following sections:
  - “Creating new views” on page 17
  - “Configuring existing views” on page 18
  - “Deleting views” on page 18
  - “Setting the default view” on page 19
- **Set Options link:** The Set Options link opens the Set Options dialog box (see Figure 14 on page 20), which allows you to set the following system-wide or component settings:
  - Refresh rate of the views that display current statistics
  - Frequency of captured historical data captured for each component type
  - Reset component statistics to “0”
  - Capture component state changes
  - Delete component state change log
  - Delete historical statistics for all components in the system
  - Delete business object probe data log

For instructions on using the Set Options dialog box, refer to the following sections:

- “Setting the refresh rate for runtime values” on page 20
- “Setting the frequency for capturing historical data” on page 20

- “Resetting runtime statistic values” on page 21
- “Capturing state changes” on page 22
- “Deleting the state change log” on page 22
- “Deleting historical statistics” on page 22
- “Deleting the business object probe data file” on page 22
- **Default view:** A default view is displayed in the body of the Web page when you log on to System Monitor. The first time you open the tool, the System Overview view included with the product is displayed. To change the default view displayed, see “Setting the default view” on page 19.
- **Logoff link:** The Logoff link allows you to log off System Manager.
- **Help link:** The Help link opens an HTML page with the following information:
  - A link to download the documentation set for the IBM WebSphere InterChange Server product
  - A directory location on your local machine where you can launch the Table of Contents file with links to help topics. This assumes you have already downloaded the documentation set.

## Using Views

You can either begin monitoring the system using the default views, or you can add, change, or delete views before monitoring the system. The following sections describe how to use existing views or create and configure views from System Monitor. A **view** is a Web page that contains one or more monitors. Several default views are included in the installation of System Monitor. You may use these default views or create new views. Before you can create and configure views, you must log on to System Monitor. For instructions on logging on to System Monitor, see “Logging on to System Monitor” on page 13.

This section covers the following topics:

“Using default views”

“Creating new views” on page 17

“Configuring existing views” on page 18

“Configuring existing views” on page 18

### Using default views

The default views that are shipped with System Monitor are described in Table 3. The table describes which monitor or monitors are contained in the view, as well as which display option is used. For descriptions of default monitors and samples of display options, see Table 1 on page 3 and “Sample display options” on page 6, respectively.

*Table 3. Default views*

Default view	Monitor(s) and display options
System Overview	System Overview monitor displayed in a table tree
Connector	<ul style="list-style-type: none"> <li>• Business Objects monitor displayed in a table</li> <li>• Connector Subscriptions monitor displayed in a table</li> </ul>
Collaboration Overview	Collaboration Statistics monitor displayed in a table

Table 3. Default views (continued)

Default view	Monitor(s) and display options
Collaboration	<ul style="list-style-type: none"> <li>• Collaboration Events monitor displayed in bar chart, and</li> <li>• Event Rate monitor displayed in a meter</li> </ul>
Collaboration History	<ul style="list-style-type: none"> <li>• Historical Collaboration Flows monitor displayed in a bar chart</li> <li>• Historical Collaboration Flows monitor displayed in a line chart</li> </ul>
Connector Overview	Connector Statistics monitor displayed in a table
Maps and Relationships	<ul style="list-style-type: none"> <li>• Map Status monitor displayed in a table</li> <li>• Relationship Status monitor displayed in a table</li> </ul>
Server Statistics	<ul style="list-style-type: none"> <li>• Historical Server Statistics displayed in a stacked bar chart</li> <li>• Database Connections displayed in a table</li> <li>• Message Queues displayed in a table</li> </ul>
Server History	<ul style="list-style-type: none"> <li>• Historical Server Statistics displayed in a bar chart</li> <li>• Historical Server Flows displayed in a line chart</li> </ul>
Flow Control	Flow Control monitor displayed in a table
State Change Log	State Change Log monitor displayed in a table

## Creating new views

The following instructions describe how to create a new view:

1. Click Create and Configure Views in the left frame of System Monitor. The Create and Configure Views dialog box appears (see Figure 12).

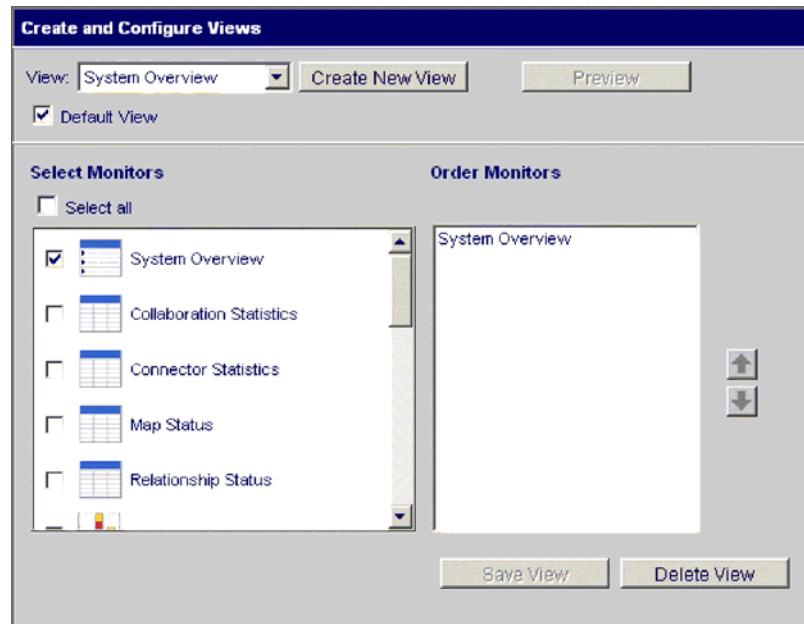


Figure 12. System Monitor, Create and Configure Views dialog box

2. Click the Create New View button. The View Name dialog box appears.

3. Type a name for the view in the View Name field, then click OK. The new view name appears in the View field of the Create and Configure Views dialog box.
4. Select one or more monitors from the Select Monitor(s) group box, or choose "Select all" to select all the monitors listed. Your selections appear in the Order Monitors group box.
5. Use the up and down arrows to the right of the Order Monitors group box to put the monitors in the order you want to view them, from top to bottom.
6. Click Preview if you want to see a preview of the new view.
7. Click Save View. A "View was saved successfully" message appears. Also, the new view appears immediately in the left frame of System Monitor under Views.

### **Configuring existing views**

The following instructions describe how to make changes to an existing view:

1. Click Create and Configure Views from the left frame of System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 17).
2. Select the view you want to change from the View drop-down menu.
3. Add monitors to or remove monitors from the view, from the Select Monitors group box. The revised monitors for the view appear in the Order Monitors group box.
4. Use the up and down arrows to the right of the Order Monitors group box to put the monitors in the order you want to view them, from top to bottom.
5. Click Preview if you want to see a preview of the new view.
6. Click Save View. A "View was saved successfully" message appears.

### **Deleting views**

The following instructions describe how to delete a view:

1. Click Create and Configure Views from the left frame of System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 17).
2. Select the view you want to delete from the View drop-down menu.
3. Click Delete View. A message appears, asking if you are sure you want to delete the view.
4. Click OK. The view is removed from the Views list in the left frame of System Monitor.

## **Fine-tuning System Monitor**

You can make adjustments to many of the elements of System Monitor, fine-tuning the level of system data you can monitor. These adjustments are described in the following sections:

"Setting the default view" on page 19

"Customizing the visual appearance of the monitors" on page 19

"Setting the refresh rate for runtime values" on page 20

"Setting the frequency for capturing historical data" on page 20

"Resetting runtime statistic values" on page 21

"Capturing state changes" on page 22

“Deleting the state change log” on page 22

“Deleting historical statistics” on page 22

“Deleting the business object probe data file” on page 22

### Setting the default view

The default view is the view you first see when you log on to System Monitor. To change the default view, do the following:

1. Click Create and Configure Views from the left frame of System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 17).
2. Select the view you want to be the default view from the View drop-down menu.
3. Select the Default View check box.
4. Click Save View. A “View was saved successfully” message appears. The next time you log on to System Manager, the view you selected as the default view will be displayed.

### Customizing the visual appearance of the monitors

The display options of monitors can be customized by changing the preferences of the display options. To change the appearance of a monitor, do the following:

1. While viewing a monitor, click the chart icon in the upper right corner. The Preferences dialog box appears for that particular display option in that monitor. Figure 13 is an example of the Table Preferences dialog box.

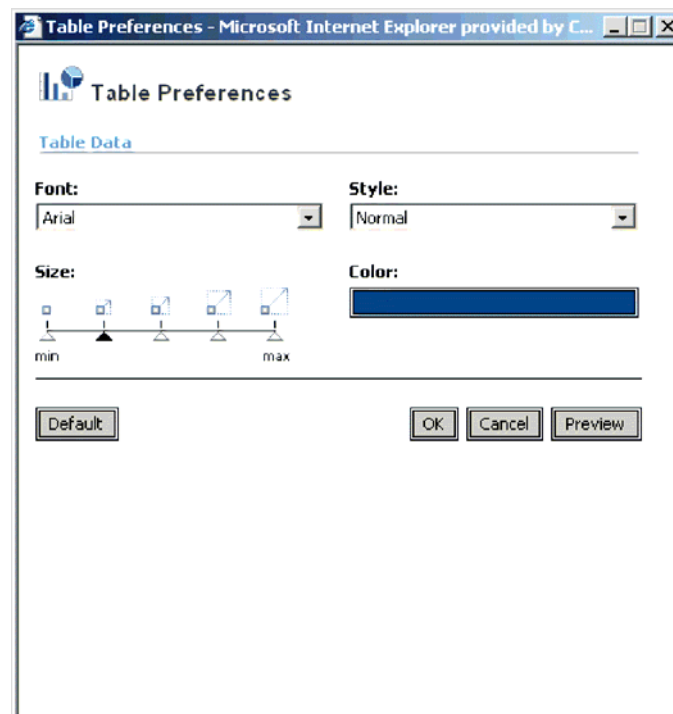


Figure 13. System Monitor, Table Preferences dialog box

2. From the Preferences dialog box, select the appearance options you want to change. For a list of what appearance options are available with each display option, see the Runtime configurable properties column of Table 2 on page 6.
3. Click Preview to see a preview of the changes you made.

4. Click OK. The changes you made appear in the monitor.

**Note:** When you change the preferences of a display option, the changes appear in all monitors that use that particular display option.

**Note:** If you want to return the monitor to its original appearance, open the Preferences dialog box, select Default, then click OK.

## Setting the refresh rate for runtime values

Some monitors display runtime values of a component. For these monitors, you can specify how often you want statistics to be refreshed.

**Note:** The refresh rate you set is for the system as a whole, not for individual components.

To set the refresh rate for monitored runtime values, do the following:

1. Click Set Options from the left frame of System Monitor. The Set Options dialog box appears (see Figure 14).

The screenshot shows the 'Set Options' dialog box with the following details:

- How often do you want to update statistics?**
  - Refresh Rate: 0 seconds
  - Submit button
- How frequently should historical data be captured?**
  - Component Type: Server
  - Frequency: NONE
  - Review all interval settings link
  - Submit button
- Do you want to reset component statistics?**
  - Component Type: Server (all components)
  - Submit button
- Do you want to capture state changes of a particular component?**
  - Component Type: Server
  - Capture State Changes checkbox
  - Submit button
- Do you want to delete the state change log for all components?**
  - Start Date: [calendar icon]
  - End Date: [calendar icon]
  - Delete button
- Do you want to delete the historical statistics for all components?**
  - Start Date: [calendar icon]
  - End Date: [calendar icon]
  - Delete button
- Do you want to delete the data for a Business Object Probe?**
  - Business Object Probe: JTextPT:BOProbe.LinkTransition\_46.LinkTransition\_46
  - Start Date: [calendar icon]
  - End Date: [calendar icon]
  - Delete button

Figure 14. System Monitor, Set Options dialog box

2. Enter a number in the Refresh Rate field to specify the number of seconds you want to set for the refresh rate, then click the Refresh Rate Submit button.

## Setting the frequency for capturing historical data

The rate at which historical data is captured is configurable. To set this rate, do the following:

1. Click Set Options from the left frame of System Monitor. The Set Options dialog box appears (see Figure 14).
2. In the "How frequently should historical data be captured?" section, click the "Review all interval settings" link. The Historical Statistics Interval Rates dialog box appears (see Figure 15 on page 21).

Historical Statistics Interval Rates	
<b>Server</b>	
DOCUMENTATION02DEV420Server	NONE
<b>Collaborations</b>	
CustomerSync_ClarifyToSAP	NONE
OrderStatusQuery_WebToCaesar	NONE
PassThroughTest_PortToCaesar	NONE
PurchasingPro_RequisitionPassThrough_PortToJText	NONE
<b>Connectors</b>	
CaesarConnector	NONE
ClarifyConnector	NONE
EMailConnector	NONE
JDBCConnector	NONE
JTextConnector	NONE
PortConnector	NONE
SAPConnector	NONE
SiebelConnector	NONE
XMLConnector	NONE
Submit Changes	

Figure 15. System Monitor, Historical Statistics Interval Rates dialog box

3. Set the interval rates for the server, and for each collaboration object and connector. The interval rates to choose from are as follows:
  - NONE
  - 15 minutes
  - 30 minutes
  - 1 hour
  - 4 hours
  - 12 hours
  - 24 hours
4. Click Submit Changes to submit all of the interval rates for all of the components.

**Note:** Alternatively, you can set the interval rate for a single component by *not* selecting the "Review all interval settings" link. Instead, select the component from the Component Type drop-down menu and the interval rate from the Frequency drop-down menu, then click the Submit button in that section.

### Resetting runtime statistic values

The runtime statistics are kept in memory from the time the server is started. If the server is running for several days or weeks, these values can become very large. You can reset the value of runtime statistics of a component to "0" by doing the following:

1. Click Set Options from the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 20).
2. In the "Do you want to reset component statistics?" section, select the component from the Component Type drop-down menu.
  - If you select Server, then runtime statistics for all components are reset.

- If you select Collaboration or Connector, then select the component from the Component drop-down list. Only statistics for that component are reset.
3. Click the Submit.

### **Capturing state changes**

You can capture state changes for each component and send them to a log file. To configure this, do the following:

1. Click Set Options from the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 20).
2. Under the "Do you want to capture state changes of a particular component?" section, select the component from the Component Type drop-down menu.

**Note:** If you selected Collaboration or Connector as the component type, you are prompted to select a particular collaboration object or connector.

3. Select the Capture State Changes check box, then click the Submit button in that section.

### **Deleting the state change log**

As the state change log grows, you may need to delete old data. You can delete the log for a particular time period by doing the following:

1. Click Set Options from the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 20).
2. Under the "Do you want to delete the state change log for all components?" section, click the calendar icons to enter the start date and end date for the data to be deleted, then click the Delete button in that section.

### **Deleting historical statistics**

As the historical data grows, you may need to delete old data. You can delete historical data for a particular time period by doing the following:

1. Click Set Options from the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 20).
2. Under the "Do you want to delete the historical statistics for all components?" section, click the calendar icons to enter the start date and end date for the data to be deleted, then click the Delete button in that section.

### **Deleting the business object probe data file**

As the business object probe data grows in size, you may need to delete old data. You can delete the data for a particular time period by doing the following:

1. Click Set Options from the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 20).
2. Under the "Do you want to delete the data for a business object probe?" section, select the business object probe from the Business Object Probe drop-down list, then click the calendar icons to enter the start date and end date for the data to be deleted, then click Delete.

---

## **Using the monitoring functions of InterChange Server Component Management view of System Manager**

Using the InterChange Server Component Management view of System Manager, you can monitor and get informational messages for all component status changes in the IBM WebSphere ICS system. From this tool, you can also start, stop, pause, and shut down IBM WebSphere InterChange Server Components, as well as



change component properties. For instructions on starting, stopping, and pausing components or on changing component properties, see Chapter 2, “Operating components of the system,” on page 47.

The following section describes how to connect to an InterChange Server instance and view component statistics from the InterChange Server Component Management view of System Manager.

## Connecting to an InterChange Server instance

To use the InterChange Server Component Management view of System Manager, you must first connect to an InterChange Server instance. The following steps describe how to do this.

1. From the InterChange Server Component Management view of System Manager (lower left quadrant of System Manager perspective), right-click the ICS instance you want to connect to, then select Connect (see Figure 16).




Figure 16. Connect to InterChange Server

2. Enter the following information in the Login dialog box that appears:
  - In the User Name field, type or select the user name for the server.
  - In the Password field, type the password for the server.
  - Click OK.

When a connection is made, the light on the ICS instance icon changes from red to green, and any objects that have been deployed to that server appear in folders beneath the server (see Figure 17).

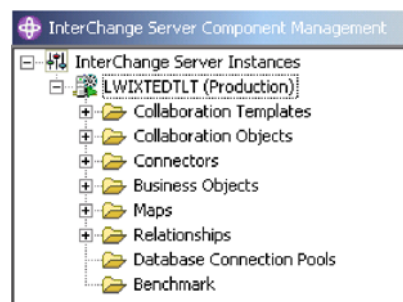


Figure 17. InterChange Server Component Management view, showing connection to ICS

## Viewing and using statistics

The InterChange Server Component Management view of System Manager allows you to monitor statistics for the IBM WebSphere ICS environment to help you better manage the IBM WebSphere ICS system. Statistics can be viewed for InterChange Server, collaboration objects, and connectors.

By watching and becoming familiar with your system's normal operating statistics, when problems occur, you can use the monitors to identify and isolate problems, and pinpoint problems in flow processing.

Monitoring your system's statistics can help you to optimally configure your system's resources. The statistics windows show currently configured parameters and provide graphs that track resources during flow processing. You can easily see if your system resources are used efficiently or if they need to be adjusted. The following topics describe the information in the Server Statistics window, the Collaboration Object Statistics window, and the Connector Statistics window:

"Server statistics"

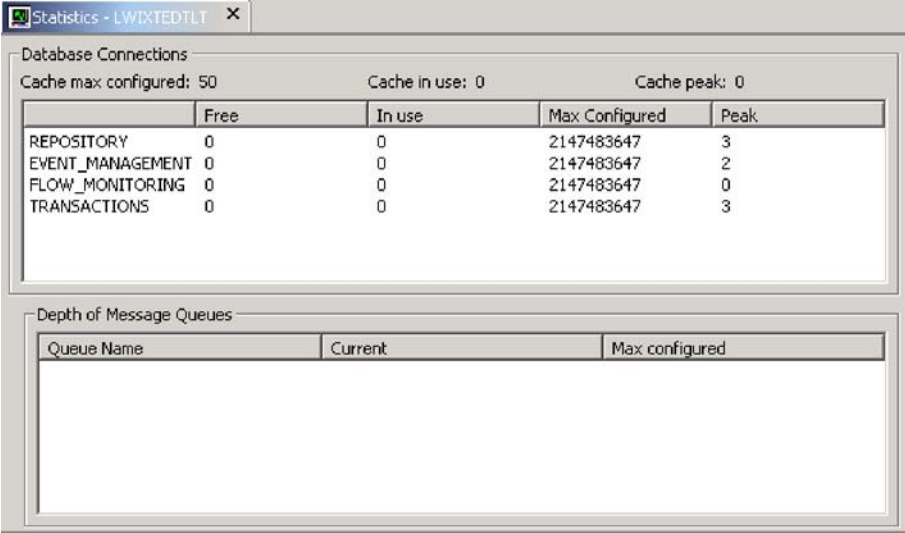
"Collaboration object statistics" on page 26

"Connector statistics" on page 27

**Note:** Before you can see any system statistics, System Manager must be connected to an ICS instance. For instructions on connecting to an ICS instance, see "Connecting to an InterChange Server instance" on page 23.

### Server statistics

To check InterChange Server statistics, right-click the ICS instance in the InterChange Server Component Management view, then select Statistics. The server statistics display in the upper-right quadrant of System Manager (see Figure 18).



The screenshot shows a window titled "Statistics - LWIXTEDTLT". It contains two main sections:

**Database Connections**

Cache max configured: 50      Cache in use: 0      Cache peak: 0

	Free	In use	Max Configured	Peak
REPOSITORY	0	0	2147483647	3
EVENT_MANAGEMENT	0	0	2147483647	2
FLOW_MONITORING	0	0	2147483647	0
TRANSACTIONS	0	0	2147483647	3

**Depth of Message Queues**

Queue Name	Current	Max configured
------------	---------	----------------

Figure 18. Statistics window for ICS

The information in the System Statistics window is covered in the following topics:

"Database connections" on page 24

"Depth of message queues" on page 25

**Database connections:** Using the Database connections section of the Statistics window for ICS, you can find out how many database connections the InterChange Server system's connection cache is currently using and the peak

amount used since the server was booted. This can help you tune InterChange Server's interaction with the underlying Database server. By using the parameters in the InterchangeSystem.cfg file and the respective underlying database server's .cfg files, you can configure the optimal number of connections.

Using this section of the System Statistics window, look for the connection pool that is consuming the most number of connections. This can help you configure InterChange Server to meet the maximum database connections constraint or increase the maximum number of connections for this pool.

The database parameters contained in the DB\_CONNECTIVITY section of the InterchangeSystem.cfg file govern the overall interactions between InterChange Server and the database management system (DBMS).

For information about these parameters, see the *System Installation Guide for UNIX or for Windows*.

The Database connections area shows statistics for:

Cache max configured	The maximum number of connections configured. This is the value the attribute MAX_CONNECTIONS. If this attribute is not configured, it displays Default.
Cache in use	The current number of connections used from the connection cache.
Cache peak	The maximum number of connections used by the server from its connection cache since the server was booted.

The area below the cache statistics lists the system and dynamic connection pools. The system pools are REPOSITORY, EVENTS\_MANAGEMENT, FLOW\_MONITORING, and TRANSACTIONS. The dynamic pool is the Relationship pool. The following details for each of these pools are maintained:

Free	The current number of available connections in the connection pool.
In use	The current number of connections used by this connection pool.
Max Configured	The maximum number of connections configured. This is the value in the InterchangeSystem.cfg file for the attribute MAX_CONNECTIONS in the respective subsections of the different connection pools (Event Management, Transactions, Repository). If this attribute is not configured, it displays Default.
Peak	The highest number of connections used by the server from this pool since the server was booted.

**Depth of message queues:** The Depth of Message Queues area of the Statistics window for ICS shows a list of all the subscription queues in the configured queue manager. The Depth of Message Queues area shows statistics for:

Queue Name	The name of the subscription queue.
Current	The number of messages currently in the queue.

This does not include subscriptions messages that are in the work-in-progress (WIP) queue.

Max Configured

The maximum number of physical messages that can exist on the queue.

## Collaboration object statistics

To check collaboration object statistics, right-click the collaboration object whose statistics you want to view, then select Statistics. The statistics for that collaboration object appear in the upper-right quadrant of System Manager (see Figure 19).

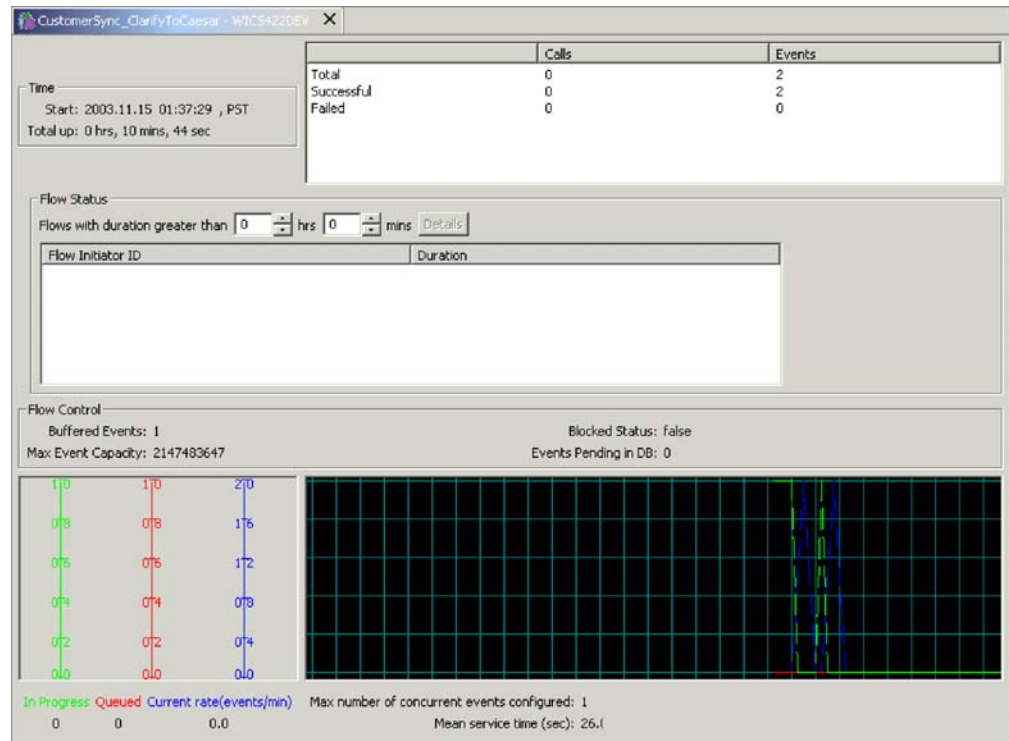


Figure 19. Statistics window for collaboration objects

**Note:** Statistics for a collaboration group, as a whole, are not maintained. Each collaboration member in a collaboration group maintains its own statistics. The statistics among group members may differ.

Check the Failed flows statistic for an increase in the normal failure rate. A failure can be caused by several situations, including the unavailability of a connector, corrupt data, and so forth. This number should be kept as low as possible, since some user intervention is needed to resubmit the failures. This count is retained when the collaboration is paused, and it is reset when the collaboration is stopped.

**Top section:** The top section of the window provides statistics about the running collaboration object. You can quickly see when the collaboration object was started, how long it has been running, the number of access calls from Web-based servlets, the number of successful and failed flows, and the total number of flows that have been processed.

The Maximum number of concurrent events reflects the maximum number of concurrent processes of event-triggered flows. For detailed information about concurrent flows, see the *System Implementation Guide*.

**Flow Status section:** Use this area to search for flows taking longer than the specified time. It can help you recognize and get details about these flows such as their FlowEventID and related application.

Enter a duration using the minutes/hours selectors to list flows whose processing time exceeds this number. The Details button provides additional information about these flows such as FlowInitiatorID, associated connector, business object, and application.

**Flow Control section:** This section displays the number of buffered events and the number of events pending in the database. It also displays two configurable Flow Control properties: Max Event Capacity and Blocked Status. Use this section to monitor the Flow Control of the collaboration object and to determine if you need to reconfigure the Flow Control properties of the collaboration object. For instructions on reconfiguring Flow Control properties, see “Configuring flow control for collaboration objects” on page 74.

**Bottom section:** Use the In Progress, Queued, and Current rate areas to monitor the number of flows that are queued, the number that are currently processing, and the rate at which the flows are processed.

Use the number of the mean service time during normal processing as a base to determine if processing rates are increasing. During normal system operation, this number should be fairly constant. A noticeable increase might reflect a problem such as a network or application slowdown or other situation that needs to be resolved.

Use the queued Events statistic to help tune the collaboration for concurrent flow processing, if necessary. If the installation consistently shows long queues, an option is to increase the number of concurrent event-triggered flows for the collaboration and restart the collaboration. Increasing the number of concurrent flows increases the system process size and may require additional database connections.

### **Connector statistics**

To check connector statistics, right-click the connector whose statistics you want to view, then select Statistics. The statistics for that connector appear in the upper-right quadrant of System Manager (see Figure 20 on page 28).

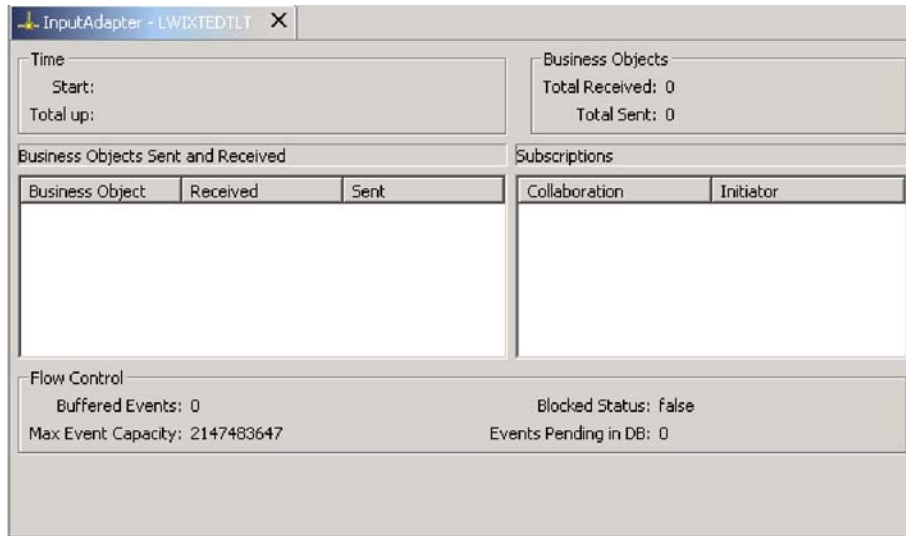


Figure 20. Statistics for connectors

The statistics window for connectors provides information about the running connector. It shows the connector's application, when the connector was started, how long it has been running, the number of business objects it has received and sent, and the Flow Control information.

**Time:** This section displays the start time of the connector and how long it has been running.

**Business objects:** This section displays the total number of business objects received and sent during the time the connector has been running.

**Business objects sent and received:** This area lists the names of the business objects the connector has sent and received. If the number of business objects sent does not match the number received, some business objects might not have been processed completely.

**Subscriptions:** This area lists the subscriptions the collaboration subscribes to and the business object name and verb for that subscription. Check the list of subscriptions to verify that the names of the collaborations and initiators are all present and that they are supposed to be there.

**Flow control:** This section displays the number of buffered events and the number of events pending in the database. It also displays two configurable Flow Control properties: Max Event Capacity and Blocked Status. Use this section to monitor the Flow Control of the connector and to determine if you need to reconfigure the Flow Control properties of the connector. For instructions on reconfiguring Flow Control properties, see "Configuring flow control for connectors" on page 68.

---

## Using the SNMP agent

The SNMP agent that can be installed with the WebSphere InterChange Server (ICS) allows a remote SNMP manager to monitor and perform limited management of InterChange Servers, collaborations, and connectors, based on a MIB (Management Information Base).

**Note:** The SNMP agent is compatible with only version-1 and version-2 SNMP managers.

The SNMP agent allows multiple InterChange Servers in an enterprise to be managed from a single agent. Conversely, several SNMP agents can manage one InterChange Server.

**Note:** WebSphere ICS does not provide protection against multiple management clients (such as System Monitor, Process Designer, or SNMP agent instances) that may be performing simultaneous management of the same component. As a result, it is possible for several clients to manage the same component and cause conflicting behavior. Use the SNMP agent with caution when other management clients are operating on the same InterChange Server.

This section covers the following topics:

“How SNMP works” on page 29

“How the SNMP agent and SNMP manager communicate” on page 30

“What SNMP manages” on page 31

“How to use the SNMP agent” on page 32

“Using the SNMP Agent Configuration Manager” on page 34

## How SNMP works

This section provides a basic overview of SNMP architecture and how the WebSphere ICS SNMP agent fits into that architecture. The following topics are covered:

“SNMP architecture”

“Management Information Base (MIB)” on page 30

“Community names” on page 30

### SNMP architecture

Simple Network Management Protocol (SNMP) is an application-layer protocol used to manage devices and processes. SNMP architecture consists of the following three components:

- SNMP manager
- SNMP agent
- Managed devices

The SNMP manager executes the applications that monitor and control managed devices. SNMP managers do not communicate directly with the managed device, but rather, they communicate through the SNMP agent. An SNMP manager is not provided as part of the ICS installation.

The SNMP agent is the entity that communicates directly with the device being managed. Its function is to receive requests from the SNMP manager, then communicate with the managed devices to process those requests. The SNMP agent is provided as an optional component of the WebSphere ICS installation. The SNMP agent can be started manually or as a Windows service.

The system devices that can be managed using the SNMP agent are InterChange Server, collaborations, and connectors. For detailed information about what types of tasks the SNMP manager can perform on each of these managed devices, see “What SNMP manages” on page 31.

### Management Information Base (MIB)

The SNMP agent uses a Management Information Base (MIB) to retrieve information about a managed device. A MIB is a collection of information that is organized hierarchically, and is like an index to the managed device. An object identifier, or object ID, uniquely identifies a managed object in the MIB hierarchy. For example, in the WebSphere ICS MIB, there is an object ID for the status of a connector. It is this object ID that is managed using SNMP.

Two MIB definitions are provided with the WebSphere ICS product: `cwsnmpagent.mib` and `cwsnmpagent_v1.mib`. These files are located in the `ProductDir\snmp` directory, where `ProductDir` is the directory where you installed the WebSphere ICS product. If your SNMP manager supports SNMP version-2, import `cwsnmpagent.mib` into your SNMP manager; if your SNMP manager supports SNMP version-1, import `cwsnmpagent_v1.mib` into your SNMP manager. You import the MIB file after you have installed and configured the SNMP agent. For instructions on installing and configuring the SNMP agent, see the *System Installation Guide for UNIX or for Windows*. For instructions on importing the MIB file to the SNMP manager, refer to the documentation provided with your SNMP management software.

**Note:** The MIB definitions for the SNMP agent are not configurable, but the SNMP agent MIB table is configurable. For instructions on configuring the SNMP agent MIB table, see “Using the SNMP Agent Configuration Manager” on page 34.

### Community names

Access control within SNMP version-2 is supported through SNMP community names. Community names function like passwords, allowing various users to manage system components by accessing the SNMP agent using a community name. For instructions on configuring community names, see “Configuring community names in the SNMP Agent Configuration Manager” on page 36.

## How the SNMP agent and SNMP manager communicate

The SNMP agent and SNMP manager communicate using the Simple Network Management Protocol. Table 4 describes the requests and notifications exchanged between the SNMP agent and SNMP manager.

Table 4. Communication between SNMP agent and SNMP manager

Request or notification	Description
Get	The SNMP manager sends this request to the SNMP agent to get information about the device or one of its managed components.
GetNext	The SNMP manager sends this request to the SNMP agent to get information about the component next to the one requested previously. This is used to iterate through a table of components.
GetBulk	The SNMP manager sends this request to the SNMP agent to get an entire table of data.
Set	The SNMP manager sends this request to the SNMP agent to set a configurable parameter in the managed devices. It is also used to start and stop components.



Table 4. Communication between SNMP agent and SNMP manager (continued)

Request or notification	Description
Trap	A trap is an asynchronous notification sent by the SNMP agent to the SNMP manager when the status of a component in the managed device changes, and the SNMP manager has expressed interest in such status changes.

## What SNMP manages

The following sections list the SNMP operations that can be performed on the managed devices by the SNMP manager. For the most up-to-date operations, refer to the current MIB file, `cwsnmpagent.mib` (for SNMP version-2) or `cwsnmpagent_v1.mib` (for SNMP version-1), located in the `ProductDir\snmp` directory. This section covers the following topics:

“SNMP management of InterChange Server” on page 31

“SNMP management of collaboration objects” on page 31

“SNMP management of connectors” on page 32

### SNMP management of InterChange Server

The SNMP manager can perform the following operations on InterChange Server through the SNMP agent:

- Register the SNMP manager to monitor the server, and unregister it.
- Query a server’s status (that is, running or stopped)
- Receive traps when a server’s status changes
- Query for the following server parameters:
  - Time of the last boot
  - Amount of time the server has been operating (uptime)
  - Number of events processed since boot time
  - Number of failed events since boot time
  - Available memory in the server (total/free)

**Note:** To access InterChange Server, the SNMP manager must first send a Set request to the SNMP agent to register the manager’s interest in monitoring a specific InterChange Server. As part of the request, the manager sends a password (previously provided to the system administrator), in addition to the standard parameters such as the community name and name of the machine where InterChange Server is installed.

### SNMP management of collaboration objects

The SNMP manager can perform the following operations on collaboration objects through the SNMP agent:

- Query a collaboration object’s status (that is, running or stopped) and mode (recovery, normal, or in-doubt)
- Receive traps when a collaboration object’s status changes
- Start or stop a collaboration object
- Query for the following collaboration object parameters:
  - Time the collaboration object was started
  - Amount of time the collaboration object has been running (uptime)

- Number of successful flows (event and access)
- Number of failed flows (event and access)
- Number of flows processed (event and access)
- Number of flows queued

### **SNMP management of connectors**

The SNMP manager can perform the following operations on connectors through the SNMP agent.

- Query a connector's status (that is, running or stopped)
- Receive traps when a connector's status changes
- Start or stop a connector
- Query for the following connector parameters:
  - Time the connector was started
  - Amount of time the connector has been running (uptime)
  - Application name
  - Number of service call request operations completed
  - Number of events retrieved

## **How to use the SNMP agent**

This section describes how to use the SNMP agent to monitor the ICS system. Before you can use the SNMP agent, you must do the following:

- Ensure that the SNMP agent is installed and configured. For instructions on installing and configuring the SNMP agent, see the *System Installation Guide for UNIX or for Windows*.
- After the SNMP agent is installed and configured, you must import the correct MIB file into your SNMP manager. For SNMP managers that support SNMP version-2, import `cwsnmpagent.mib`, and for SNMP managers that support SNMP version-1, import `cwsnmpagent_v1.mib`. These files are located in the `ProductDir\snmp` directory.

**Note:** This document does not cover instructions on how to use an SNMP manager. For instructions on how to use your SNMP manager, including instructions on how to import a MIB file, refer to the documentation that came with your SNMP management software.

This section covers the following topics:

“Starting the SNMP agent” on page 32

“Stopping the SNMP agent” on page 33

“Reconfiguring the SNMP agent” on page 33

### **Starting the SNMP agent**

To start the SNMP agent:

#### UNIX

To start, stop, or get status on the SNMP agent on UNIX, run the `snmpagent_manager` script.

If you run the SNMP agent on a UNIX operating system and the SNMP agent is configured to the default port (1161) or to any port number less than 1024, you must be root to run the SNMP startup script. If the SNMP agent is configured to a port number greater than or equal to 1024, a non-root user can start the script.

#### Windows

Click Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration SNMP > SNMP Agent.

### Stopping the SNMP agent

To stop the SNMP agent:

#### UNIX

Use the `snmpagent_manager` script.

#### Windows

Stop the SNMP agent from the Windows Services control panel.

### Reconfiguring the SNMP agent

Configuration information for the SNMP agent is stored in a configuration file named `cwsnmpagent.cfg` in the `ProductDir\snmp\config` directory.

To change the default values, edit the file as necessary.

The values contained in the file are specified as:

*ParameterName: value*

Table 5 lists the parameters used for the operation of the SNMP agent.

Table 5. SNMP agent configuration file parameters

Parameter	Description	Values	Default
TraceLevel	Defines the verbosity of the trace information.	0-5	0
LogFile	Path to the log file.		<code>cwsnmpagent.log</code> , located in the <code>ProductDir\snmp\log</code> directory
AgentStateFile	Path to the file that contains the agent's state.		<code>cwsnmpagent.sts</code> , located in the <code>ProductDir\snmp\state</code> directory.

Table 5. SNMP agent configuration file parameters (continued)

Parameter	Description	Values	Default
PollInterval	Not all information required by the SNMP agent is available through callbacks, and certain information needs to be obtained through periodic polling. This parameter specifies the polling interval, in seconds	0 (no polling) and up	30
Port	The port on which the SNMP agent listens for requests from SNMP managers	A valid port number	1161 (the default SNMP port number for UNIX) 161 (the default SNMP port number for Windows)

If you want to make changes to the SNMP configuration, the SNMP Configuration wizard provides fields for the information in Table 5. The wizard creates (or modifies) a `cwsnmpagent.cfg` file based on the values in these fields.

After changing the configuration file, shut down the SNMP agent and restart it.

## Using the SNMP Agent Configuration Manager

The SNMP Agent Configuration Manager allows you to configure the MIB tables associated with the SNMP agent. For more information about MIB tables, see “Management Information Base (MIB)” on page 30. These tables include the Community table, the Trap Forwarding table, and the Server Access table.

This tool is required when a third-party SNMP manager is not able to create a new MIB table entry for the SNMP agent. Most SNMP managers already have this functionality built in, but some do not. If you use an SNMP manager that does not allow you to configure MIB tables, you must use the SNMP Agent Configuration Manager to configure the MIB table associated with the SNMP agent. Even if your SNMP manager does have MIB table configuration capabilities, this tool is recommended for configuring the WebSphere ICS SNMP agent.

**Note:** The WebSphere ICS SNMP solution is metadata driven; therefore, you need define only the servers to monitor, not the specific connectors or collaborations.

This section covers the following topics:

“Installing the SNMP Agent Configuration Manager” on page 35

“Starting the SNMP Agent Configuration Manager” on page 35

“Connecting to the SNMP agent” on page 35

“Configuring community names in the SNMP Agent Configuration Manager” on page 36

“Configuring trap forwarding entries in the SNMP Agent Configuration Manager” on page 37

“Configuring server access entries in the SNMP Agent Configuration Manager” on page 38

## Installing the SNMP Agent Configuration Manager

When you install the SNMP agent, the installer automatically installs the SNMP Agent Configuration Manager. For instructions on installing the SNMP agent, refer to the *System Installation Guide for UNIX* or *for Windows*.

## Starting the SNMP Agent Configuration Manager

To start the SNMP Agent Configuration Manager, do one of the following tasks, depending on your operating system:

### UNIX

Run the `start_snmpconfig.bat` script located in the *ProductDir/bin* directory.

### Windows

Click Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration SNMP > SNMP Agent Configuration Manager.

The SNMP Agent Configuration Manager appears.

Before you can begin editing the MIB table, you must connect to the SNMP agent. See “Connecting to the SNMP agent” for instructions.

**Note:** The SNMP Agent Configuration Manager must be run on the same machine as the SNMP agent.

## Connecting to the SNMP agent

To connect the SNMP Agent Configuration Manager to a running SNMP agent, do the following:

1. Start the SNMP Agent Configuration Manager. For instructions, see “Starting the SNMP Agent Configuration Manager.”
2. Enter information in the following fields of the SNMP Agent Configuration Manager window:
  - **Agent Host:** Enter the host name or IP address of the machine where the SNMP agent is running.
  - **Port:** Enter the port number that the SNMP agent will use to listen for SNMP commands. If you do not know your SNMP agent port number, you can find it in the `cwsnmpagent.cfg`, located in the *ProductDir\snmp\config* directory, where *ProductDir* is the directory where you installed the WebSphere InterChange Server product. The default port number is 1161 for UNIX and 161 for Windows.
  - **Community:** Enter the read-write community name of the SNMP agent. By default, the read-write community name is “administrator” and the read-only community name is “public.” The difference between the two types of community names is as follows:
    - **Read-write:** Read-write access permits the user to edit MIB table components, query values, start and stop components, and register for traps.
    - **Read-only:** Read-only access permits the user to perform “Get” operations only. This user can view but not edit the MIB table components and cannot change the status of any component.

3. Click Connect. When the SNMP Agent Configuration Manager connects to the SNMP agent, the Agent Host, Port, and Community fields become disabled, and the Connect button changes to a Disconnect button.

## Configuring community names in the SNMP Agent Configuration Manager

The Community tab lists the communities that exist for the connected SNMP agent. The table has three entries:

- **Community Name:** The name of the community
- **Access:** The access right of the community (read-write or read-only)
- **Row Status:** The status of the community (active or not in service)

**Adding community names:** You add community names to the MIB table when you want to give new users permission to manage system components. To add a community name, do the following:

1. From the Community tab, click Add. The Community Table Item dialog box appears.
2. Type a community name in the Community Name field.

**Note:** Each community name must be unique in the network.

3. Select the type of access for the new community from the Access drop-down menu. The choices are "read-write" or "read-only."
4. The Row Status drop-down menu is not configurable when adding a community name. The default setting is "create & go."
5. Click OK. The new community name appears in the Community table with the Row Status set to Active.

**Editing community names:** You edit community names when you want to change the type of access or row status of a registered community name. To edit a community name, do the following:

1. From the Community tab, select the community name from the Community Name column, then click Edit. The Community Table Item dialog box appears.

**Note:** You cannot edit the community name.

2. Change the type of access from the Access drop-down menu. The choices are "read-write" or "read-only."
3. Change the status by selecting one of the following options from the Row Status drop-down menu:
  - **active:** This activates the community name, allowing it to be used to access the SNMP agent.
  - **not in service:** This deactivates the community name, but it stores the community name information in the table. Choose this option if you know you want to reactivate this community name at some point in the future.
  - **destroy:** This removes the community name entry from the MIB table.
4. Click OK. Any changes you made appear in the Community table.

**Removing community names:** You remove community names when you want to completely delete them from the MIB table. To remove a community name, do the following:

1. From the Community tab, select the community name from the Community Name column.
2. Click Remove. The community name is removed from the table.

## Configuring trap forwarding entries in the SNMP Agent Configuration Manager

A trap is an asynchronous notification sent by the SNMP agent to the SNMP manager whenever the status of a component in the managed device has changed and the SNMP manager has expressed interest in such status changes. The Trap Forwarding table has four entries:

- **Trap ID:** The unique ID given to a trap subscription request
- **Manager Host:** The host name of the SNMP manager where the notifications of a trap will be sent
- **Trap Port:** The port on which the SNMP manager listens to the trap
- **Row Status:** The status of the trap (active or not in service)

**Adding trap forwarding entries:** You add trap forwarding entries to the MIB table when you want to register a server as the recipient of information gathered by the SNMP agent. To add a trap forwarding entry to the MIB table, do the following:

1. From the Trap Forwarding tab, click Add. The Trap Forwarding Table Item dialog box appears.
2. In the Trap ID field, type an integer.

**Note:** Each Trap ID must be unique in the network.

3. In the Manager Host field, type the host name or IP address of the machine where the SNMP manager runs.

**Note:** For UNIX users who use the host name, be sure to use the correct case.

4. In the Trap Port field, type the port number that the SNMP manager uses to listen to traps.
5. The Row Status drop-down menu is not configurable when adding a trap forwarding entry. The default setting is "create & go."
6. Click OK. The new trap forwarding entry appears in the Trap Forwarding table with the Row Status set to Active.

**Editing trap forwarding entries:** You edit trap forwarding entries when you want to change the Manager Host, Trap Port, or Row Status information of registered trap forwarding entries. To edit an existing trap forwarding entry, do the following:

1. In the Trap Forwarding tab, select a Trap ID, then click Edit. The Trap Forwarding Table Item dialog box appears.

**Note:** You cannot edit the Trap ID.

2. Change the Manager Host by typing a different host name or IP address of the machine where the SNMP manager runs.

**Note:** For UNIX users who use the host name, be sure to use the correct case.

3. Change the Trap Port by typing a different port number that the SNMP manager uses to listen to traps.
4. Change the Row Status of the trap forwarding entry by selecting one of the following options from the Row Status drop-down menu:
  - **active:** This activates the trap forwarding entry, allowing it to be used to access the SNMP agent.
  - **not in service:** This deactivates the trap forwarding entry, but it stores the trap forwarding information in the table. Choose this option if you know you want to reactivate this trap forwarding entry at some point in the future.

- **destroy:** This removes the chosen community name entry from the MIB table.
5. Click OK. Any changes you made appear in the Trap Forwarding table.

**Removing trap forwarding entries:** You remove trap forwarding entries when you want to completely remove them from the MIB table. To remove a trap forwarding entry, do the following:

1. From the Trap Forwarding tab, select the Trap ID you want to remove.
2. Click Remove. The trap forwarding entry is removed from the Trap Forwarding table.

## Configuring server access entries in the SNMP Agent Configuration Manager

The server access entries allow you to link specific SNMP managers with specific InterChange Servers to be managed. The table has three entries:

- **Manager Host:** The host name of the SNMP manager
- **WebSphere InterChange Server:** The host name of the machine where InterChange Server is installed
- **Row Status:** The link status (active or not in service)

**Adding server access entries:** To create a new server access, do the following:

1. From the Server Access tab, click Add. The Server Access Table Item dialog box appears.
2. In the Manager Host field, type the host name or IP address of the machine where the SNMP manager runs.

**Note:** For UNIX users who use the host name, be sure to use the correct case.

3. In the WebSphere InterChange Server field, type the InterChange Server name or IP address.

**Note:** For UNIX users who use the host name, be sure to use the correct case.

4. The Row Status drop-down menu is not configurable when adding a Server Access entry. It is set to "create & go."
5. Click OK. The new server access entry is added to the Server Access table, with the Row Status set to "not in service."

**Editing server access entries:** You edit the server access entries of the MIB table when you want to change the row status of the server access entries. To edit a server access entry, do the following:

1. From the Server Access tab, select a Manager Host, then click Edit.
2. Change the status, select one of the following options from the Row Status drop-down menu:
  - **active:** This activates the chosen manager host, allowing it to be used to used with the SNMP agent.

**Note:** In order to change the status from "not in service" to "active," both the SNMP manager and the InterChange Server being managed must be running.

- **not in service:** This deactivates the manager host, but it stores the manager host information in the table. Choose this option if you know you want to reactivate this manager host at some point in the future.
  - **destroy:** This removes the chosen manager host from the MIB table.
3. Click OK. Any changes you made appear in the Server Access table.



**Removing server access entries:** You remove server access entries when you want to completely delete them from the MIB table. To remove a server access entry, do the following:

1. From the Server Access tab, select the server access entry you want to remove.
2. Click Remove. The Server Access entry is removed from the Server Access table.

---

## Using persistent monitoring

Persistent monitoring is a subsystem of ICS that monitors and stores historical state and statistical information of collaboration objects, connectors and the system as a whole. Prior to release 4.2.0, only current state and statistical information about system components was available for troubleshooting or maintaining the system.

This section describes how to configure system-wide and component-level persistent monitoring and how to access the results from persistent monitoring.

**Note:** You must consider the database volume requirements as well as a data deletion strategy when planning the number of components being monitored and the frequency at which they are monitored. For more information about implementing database volume requirements, see the *System Implementation Guide*.

This section covers the following topics:

“Configuring persistent monitoring”

“Accessing the results from persistent monitoring” on page 40

## Configuring persistent monitoring

You configure the various levels of persistent monitoring for system components from System Monitor. Refer to the following sections for instructions:

“Setting the refresh rate for runtime values” on page 20

“Setting the frequency for capturing historical data” on page 20

“Resetting runtime statistic values” on page 21

“Capturing state changes” on page 22

“Deleting the state change log” on page 22

“Deleting historical statistics” on page 22

“Deleting the business object probe data file” on page 22

You can also configure certain system-wide parameters of persistent monitoring from System Monitor. The following instructions explain how to do this.

1. Open the Edit Configuration tool by right-clicking the server from the Server Instances section of System Manager, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the InterchangeServer.cfg file.

2. Select the Misc tab (see Figure 21).

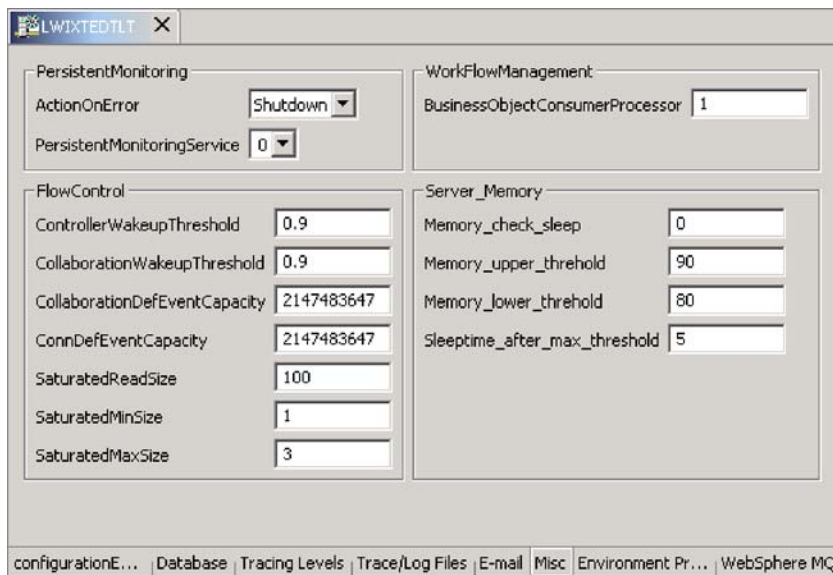


Figure 21. Edit Configuration screen, Misc tab

3. Do the following in the "Persistent Monitoring" pane:
  - a. Select **Continue** in the "Action on error" drop-down menu if you want InterChange Server to continue running in the event of errors experienced by the persistent monitoring system.  
Select **Shutdown** in the "Action on error" drop-down menu if you want InterChange Server to shut down in response to errors with the subsystem.
  - b. Select the desired tracing level in the "Persistent monitoring service" drop-down menu to specify the tracing level for the subsystem.

## Accessing the results from persistent monitoring

You access the results of persistent monitoring by using System Monitor. Two of the default views provided with System Monitor display historical state and statistical information: Collaboration History and Server History. For more information on using default views, see "Using default views" on page 16. Alternatively, you can create your own views that can contain historical data. For more information on creating views, see "Creating new views" on page 17.

---

## Using WebSphere Business Integration Monitor

WebSphere Business Integration (WBI) Monitor is a standalone monitor that allows you to monitor a variety of servers, one of which is InterChange Server (ICS). WBI Monitor is part of the WBI Modeler & Monitor product. In order for WBI Monitor to have the capability to monitor flows in ICS, the following requirements must be met:

- WBI Monitor, version 4.2.4, Fix pack 1, must be installed on the machine from where you intend to do the monitoring.
- MQ Workflow adapter must be installed and enabled with **flow monitoring**. Flow monitoring is a service that provides a view of the event information associated with each flow as it passes through a collaboration in ICS.
- InterChange Server, version 4.2.2.2, must be configured for flow monitoring.

- Enable flow monitoring for the collaboration with the Auditing Collaboration Utility provided by the WBI Monitor. For further information, refer to the Administration Guide in the WBI Monitor documentation.

For WBI Monitor installation and deployment instructions, refer to the Deployment guide in the WBI Monitor documentation. For MQ Workflow adapter installation instructions, refer to the *Adapter for MQ Workflow User Guide*.

**Important:** WBI Monitor will not work with ICS if ICS is using Microsoft SQL Server as its database.

This section covers the following topics:

“How WBI Monitor works with ICS”

“Enabling flow monitoring in the MQWorkflow adapter”

“Configuring InterChange Server for flow monitoring”

“Configuring tracing for flow monitoring” on page 46

**Note:** This guide does not cover instructions for administering WBI Monitor. For WBI Monitor administration instructions, refer to the Administration Guide in the WBI Monitor documentation.

## How WBI Monitor works with ICS

WBI Monitor is able to monitor an ICS system only when that ICS system is configured with flow monitoring. Flow monitoring is a service that provides a view of the event information associated with each flow as it passes through a collaboration in ICS.

**Note:** WBI Monitor monitors ICS events only, not the business data associated with those events.

## Enabling flow monitoring in the MQWorkflow adapter

WBI Monitor can monitor an ICS system only when the MQWorkflow adapter is enabled for flow monitoring. Flow monitoring is a service that provides a view of the event information associated with each flow as it passes through a collaboration in ICS. The following steps describe how to enable the MQWorkflow adapter with flow monitoring.

1. From System Manager, double-click the MQWorkflow connector. This opens Connector Configurator.
2. In the Standard Properties tab, scroll to the `EnableOidForFlowMonitoring` property.
3. In the Value column, select “true.”
4. Click File > Save > To Project.
5. Deploy the repository to InterChange Server.
6. Restart InterChange Server.

## Configuring InterChange Server for flow monitoring

WBI Monitor can monitor an InterChange Server (ICS) system only when that ICS system is first configured with flow monitoring. Flow monitoring is a service that

provides a view of the event information associated with each flow as it passes through a collaboration in ICS. The following sections describe how to configure ICS for flow monitoring.

“Configuring the ICS database for flow monitoring”

“Configuring maximum queue depth for flow monitoring” on page 45

“Configuring tracing for flow monitoring” on page 46

### **Configuring the ICS database for flow monitoring**

The ICS database can be configured for flow monitoring only after the ICS installation is complete. Two tools allow you to configure the ICS database for flow monitoring: the ICS Configuration Wizard and System Manager. The following steps describe how to configure the ICS database for flow monitoring, using either of these tools.

**Important:** Flow monitoring cannot be configured on ICS if Microsoft SQL Server is the ICS database.

1. Depending on which tool you use, open the tool using one of the following instructions:

#### **InterChange Server Configuration Wizard**

Click Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere InterChange Server > IBM WebSphere InterChange Server Configuration Wizard.

The InterChange Server Configuration Wizard opens.

#### **System Manager**

From the InterChange Server Component Management view of System Monitor, right-click the ICS instance whose database you want to set up with the flow monitoring service, then select Edit Configuration.

The upper-right section of System Manager becomes a tool from which you can edit the InterchangeSystem.cfg file.

2. Select the Database tab (of either the ICS Configuration Wizard or the Edit Configuration window of System Manager), then scroll to the bottom of the window to see the Flow Monitoring section (see Figure 22 on page 43 for the ICS Configuration Wizard or Figure 23 on page 44 for the Edit Configuration window of System Manager).

**Note:** If you are using Microsoft SQL Server as your ICS database, the Flow Monitoring section will be greyed out. Flow monitoring cannot be configured on ICS if Microsoft SQL Server is your ICS database.

**InterChange Server Configuration**

InterChange Server | WebSphere MQ | Database

**Database Connectivity**

Database Driver: DB2 Max Connections: 50  Unlimited  
 Max Pools: 10  
 Idle Timeout: 2

**Event Management**

Host Name: \_\_\_\_\_ Login: wicsadmin  
 Database: wicsrepos Password: \*\*\*\*\*   
 Max Connections: \_\_\_\_\_  Unlimited Port Number: \_\_\_\_\_

**Transactions**

Host Name: \_\_\_\_\_ Login: wicsadmin  
 Database: wicsrepos Password: \*\*\*\*\*   
 Max Connections: \_\_\_\_\_  Unlimited Port Number: \_\_\_\_\_

**Repository**

Host Name: \_\_\_\_\_ Login: wicsadmin  
 Database: wicsrepos Password: \*\*\*\*\*   
 Max Connections: \_\_\_\_\_  Unlimited Port Number: \_\_\_\_\_

**Flow Monitoring**

Host Name: \_\_\_\_\_ Login: wicsadmin  
 Database: wicsrepos Password: \*\*\*\*\*   
 Schema Name: mySchema Port Number: \_\_\_\_\_  
 Max Connections: \_\_\_\_\_  Unlimited

Figure 22. Database tab of the ICS Configuration Wizard

Figure 23. Database tab of the Edit Configuration window of System Manager

3. Enter information in the Flow Monitoring fields, using the information in Table 6 as a guide.

Table 6. Flow Monitoring configuration information

Configuration information	Description
Host Name	Host name of the machine where the database server resides
Database	Name of the database you created on your database server
Schema Name	Name of the database schema where the flow monitoring event tables reside. If you configure a custom schema name, you must grant CREATE, DELETE, and INSERT permission to the login identified by the schema name. For DB2 databases, the schema name can be arbitrary. For Oracle databases, the schema name is the same as the user who creates the table. The default schema name is the same as the login user name. <b>Note 1:</b> Valid values for the schema name field can contain up to 30 characters from the US-ASCII character set. The name must begin with a letter from A through Z and the first three characters cannot be SYS. Other characters in the name can include the letters A through Z and numbers 0 through 9. <b>Note 2:</b> The "Schema name" field is available only on 4.2.2.2 versions of ICS.
Max connections	Maximum number of simultaneous connections to the database
Login	Database login name

Table 6. Flow Monitoring configuration information (continued)

Configuration information	Description
Password	Database password
Port number	Port number of your database server

### Configuring maximum queue depth for flow monitoring

Maximum queue depth is a parameter that controls the maximum number of events allowed in the ICS memory before collaborations wait to enqueue additional events. During flow monitoring, many events are recorded for every flow in a traced collaboration. This results in large amounts of database activity, which can degrade performance. To prevent excessive performance degradation, the collaboration flows write monitor events to an in-memory ICS queue, then these flows are transferred to the database. In order to prevent the ICS system from running out of memory, this queue depth can be limited by configuring the maximum queue depth parameter.

The following steps describe how to configure the maximum queue depth from System Manager.

**Note:** The maximum queue depth parameter is configurable system-wide only, not on a per collaboration basis.

1. From the InterChange Component Management view, right-click the ICS instance whose maximum queue depth you want to configure, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the InterchangeSystem.cfg file.
2. Click the Misc tab (see Figure 24).

The screenshot shows the 'Misc' tab of the ICS Edit Configuration window. It contains several configuration panels:

- Persistent Monitoring:** Action on error is set to 'Shutdown', and Persistent monitoring service is set to '0'.
- WorkFlow Management:** Business object consumer processor is set to '1'.
- Flow Control:** Controller wakeup threshold, Collaboration wakeup threshold, Collaboration def event capacity, Conn def event capacity, Saturated read size, Saturated min size, and Saturated max size are all set to 0.9, 0.9, 2147483647, 2147483647, 100, 1, and 3 respectively.
- Server Memory:** Memory checker sleep time (minutes) is 0, Memory upper threshold percent is 90, Memory lower threshold percent is 80, and Connector pause time at threshold (minutes) is 5.
- Flow Monitoring:** Max flow queue depth is 500, and the Max value checkbox is unchecked.

Figure 24. Misc tab of ICS Edit Configuration window

3. In the "Max flow queue depth" field of the Flow Monitoring section, enter a number that represents the maximum number of events you want to allow in the ICS memory before collaborations wait to enqueue additional events. The default is 500. Alternatively, you can select the Max value check box. The maximum number of events allowed is 2147483647.

The changes are saved immediately to the InterchangeSystem.cfg file, but do not take effect until the server is restarted.

## Configuring tracing for flow monitoring

When you configure flow monitoring for an ICS system, you can set a desired trace level for traced events. The following steps describe how to set the trace level for flow monitoring.

**Note:** Setting the trace level for flow monitoring can be done only on a system-wide basis, not on a per collaboration basis.

1. From the InterChange Component Management view, right-click the ICS instance whose flow monitoring trace level you want to configure, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the InterchangeSystem.cfg file.
2. Select the Tracing Levels tab.
3. In the IBM WebSphere Business Integration System Trace Levels group box, look for the "Trace level" field for Flow Monitoring.
4. In the Flow Monitoring Trace level field, select a number from 0-5 to represent the desired trace level. Table 7 describes the different trace levels for flow monitoring.

Table 7. Flow monitoring trace levels

Trace level	Description
0	No tracing
1	At boot-time, display the configuration to the database
2	Not used
3	Show event enqueue/dequeue in-memory of the server
4	Show event data details such as type or sequenceNumber after in-memory queue operation
5	Add database writes



---

## Chapter 2. Operating components of the system

This chapter describes some of the tasks you may need to perform while operating an IBM WebSphere ICS system. For instructions on starting InterChange Server for the first time, refer to the *Installation Guide for UNIX or for Windows*. This chapter covers the following topics:

“Overview of operating the system” on page 47

“Operating InterChange Server” on page 48

“Using System Manager” on page 52

“Operating connectors” on page 56

“Operating collaboration objects” on page 69

“Operating maps” on page 76

“Operating relationships” on page 78

“Using Relationship Manager” on page 82

“Backing up IBM WebSphere InterChange Server components” on page 96

“Using Repos\_Copy” on page 99

“Scheduling jobs” on page 110

---

### Overview of operating the system

For IBM WebSphere InterChange Server (ICS) system, the order you must follow is to start WebSphere MQ Listener, InterChange Server, then IBM WebSphere System Manager. The connectors are automatically started when you start InterChange Server.

When shutting down an instance of InterChange Server, you have two choices. You can shut down the system gracefully or immediately. A graceful shutdown allows the system to complete work that is in progress before shutting down, whereas an immediate shutdown stops the system without allowing pending events to process.

Before you can start the IBM WebSphere ICS system, make sure all of the necessary third-party software is running. This includes the database on which the IBM WebSphere ICS repository resides. This section assumes that you have already started the system and loaded the repository. If you are starting the system for the first time, refer to the *System Installation Guide for UNIX or for Windows*.

When starting up the system, all connectors and collaborations come up in the state they were in at the time they were last shut down. For example, if a collaboration was paused at shutdown, it is paused when the system is started again.

The following tasks describe how to start the system:

1. Check that all necessary third-party software is running.
2. Start InterChange Server. See “Starting InterChange Server” on page 48.
3. Start any connectors not automatically started. See “Starting, stopping, and pausing connectors” on page 58.
4. Start collaborations. See “Collaboration object states” on page 69.
5. Start System Manager. See “Using System Manager” on page 52.
6. Start the controller gateway (if Secure Connect is installed). See the *System Implementation Guide* for more information about the Secure Connect.

---

## Operating InterChange Server

Operating InterChange Server may involve the following tasks:

“Starting InterChange Server” on page 48

“Shutting down InterChange Server” on page 49

“Changing the InterChange Server and database passwords” on page 50

## Starting InterChange Server

To start InterChange Server:

### UNIX

Run the `ics_manager -start` script.

**Note:** If the ICS password is changed, you must use the following script:  
`-start -uloginName -ppassword`

### Windows

Click Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere InterChange Server > IBM WebSphere InterChange Server.

At startup, InterChange Server reads the `InterchangeSystem.cfg` file and sets its properties according to the parameter values listed there. See the *System Installation Guide for UNIX or for Windows* for a list and description of the configuration parameters.

## Customizing InterChange Server startup parameters

To customize the InterChange Server startup:

### UNIX

Modify the `ics_manager` script. When running this script, you can use the following arguments to start, stop, or see the status of InterChange Server:

`-start`

`-stop`

`-status`

### Windows

Modify the InterChange Server shortcut or the `start_server.bat` file.

The parameters in table Table 8 customize the startup of InterChange Server.

Table 8. InterChange Server startup parameters

Parameter	Function
<code>-c configFile</code>	Name of the configuration file to be used during startup. The default is <code>InterchangeSystem.cfg</code> .
<code>-i</code>	Allows InterChange Server to start up and ignore all error messages.
<code>-p password</code>	Specifies the password to access InterChange Server. If you do not use this parameter, the <code>start_server</code> command uses the password in the <code>InterchangeSystem.cfg</code> file. Use with the <code>-u</code> parameter.
<code>-s serverName</code>	Specifies the name of the InterChange Server. The name is case-sensitive.
<code>-u loginName</code>	Specifies the user login name for InterChange Server. If you do not use this parameter, the <code>start_server</code> command uses the user login name in the <code>InterchangeSystem.cfg</code> file. Use with the <code>-p</code> parameter.
<code>-v</code>	Displays the version of InterChange Server, then exits.

## Shutting down InterChange Server

Shutting down InterChange Server stops all running collaborations and connectors, as well as InterChange Server itself. All connections to the database are closed and the machine's system resources used by InterChange Server are returned.

Depending on your operating system, you can shut down the server using one of the following methods:

### UNIX

Run the `ics_manager -stop` script.

**Note:** If the ICS password is changed, you must use the following script:

```
-stop-uloginName -ppassword
```

### Windows

In InterChange Server Component Management view of System Manager, right-click the ICS instance, then select Shut Down > Gracefully. Alternatively, you can select Shut Down > Immediately, which shuts down the server without cleanup. To determine which type of shutdown is best for you, refer to the following topics: "Graceful shutdown" on page 50 and "Immediate shutdown" on page 50.

**Note:** An HA system cannot be shut down using System Monitor, SNMP manager, or the InterChange Server Component Management view of System Manager). Shutting down the server using any of these tools causes the HA system to perform a failover. To shut down the HA system, use the MSCS

Administrator and use the Take Offline context menu option for all connectors and then InterChange Server. For instructions, see “Changing the status of a resource” on page 169. For more information about managing an HA system, see “Managing High-Availability (HA) systems” on page 167..

**Attention:** Refrain from using Ctrl+C to shut down InterChange Server. Doing so prevents the server from shutting down in an orderly manner.

### Graceful shutdown

Gracefully shutting down the system allows all currently processing and queued flows to complete before shutting down. This may take a long time because all flows waiting to be processed by a running collaboration must complete. Existing flows are processed by the collaborations, but no new flows are accepted.

If you choose to gracefully shut down the system, the following occurs:

- Connectors stop polling. No new events are generated.
- Collaboration objects finish their current work, then stop.

If the collaboration object is a member of a collaboration group, all collaboration objects in the group stop.

If messages from the connectors are in transit to the collaboration object when it stops, they remain in the messaging queues until the collaboration object starts.

- InterChange Server shuts down.

**Note:** This procedure cannot be used to shut down an HA system. For information on shutting down a high-availability (HA) system, see “Changing the status of a resource” on page 169..

### Immediate shutdown

Immediately stopping the system forces the system to shut down without processing any more flows. Running connectors and collaborations are stopped immediately. When the system is restarted, flows that were interrupted by the immediate shutdown are redelivered in the same processing order. If one of these flows wrote data to an application, when the flow is redelivered, it tries to duplicate the data and fails because the data already exists. If the collaboration processing the flow is transactional, a rollback occurs. If the flow is not transactional, it is moved to the resubmission queue. See “Failed flows” on page 151 for more information on submitting a flow that fails to process.

**Note:** Immediately stopping the system does not compromise the integrity of the data or the integrity of the IBM WebSphere ICS system.

Use this option when you need to quickly shut down the system. For example, you may want to reboot the system, but a collaboration has multiple events waiting to be processed. Shutting down gracefully may take too much time because the collaborations need to complete all existing work before stopping.

## Changing the InterChange Server and database passwords

Password encryption provides a measure of security for protecting the IBM WebSphere InterChange Server system and underlying databases from unauthorized user entry. The encrypted string for each of the passwords is stored in InterChange Server and is accessed by the server when the password must be decrypted. In the InterchangeSystem.cfg file, the encrypted password is placed in the PASSWORD\*= parameter.

The IBM WebSphere ICS administrator and database passwords are requested during system installation by Installer and are encrypted and stored when the system is rebooted at the completion of the installation. Thereafter, you can change the InterChange Server password or the database password in System Manager.

The InterChange Server user name and password are required during repository copy and restoration when the `repos_copy` command is used. See “Using Repos\_Copy” on page 99.

For instructions on changing the password for InterChange Server or for the database(s), refer to the following sections:

“Changing the InterChange Server password”

“Changing the database passwords”

### **Changing the InterChange Server password**

To change the password for InterChange Server:

1. Open System Manager.
2. Right-click the ICS instance in the InterChange Server Component Management view, then select Change Password. The Change InterChange Server Password dialog box appears.
3. Enter the current password in the Old Password field.
4. Enter a new password in the New Password field.
5. Reenter the new password in the Confirm Password field.
6. Click OK.

The encrypted password is stored in the `InterchangeSystem.cfg` file.

**Attention:** The InterChange Server password can be changed only by using this procedure. If you try to change the password by editing the password in the `InterchangeSystem.cfg` file, InterChange Server will not start.

### **Changing the database passwords**

The repository database passwords can be changed through System Manager once the IBM WebSphere ICS system is operating.

To change the database passwords:

1. In the InterChange Server Component Management view, right-click the ICS instance whose database password you want to change, then select Edit Configuration. The upper-right section of the window changes to an editing tool in which many system properties can be changed.
2. Click the Database tab to access the database configuration properties. The Server Property and Configuration window for database properties appears (see Figure 25 on page 52).

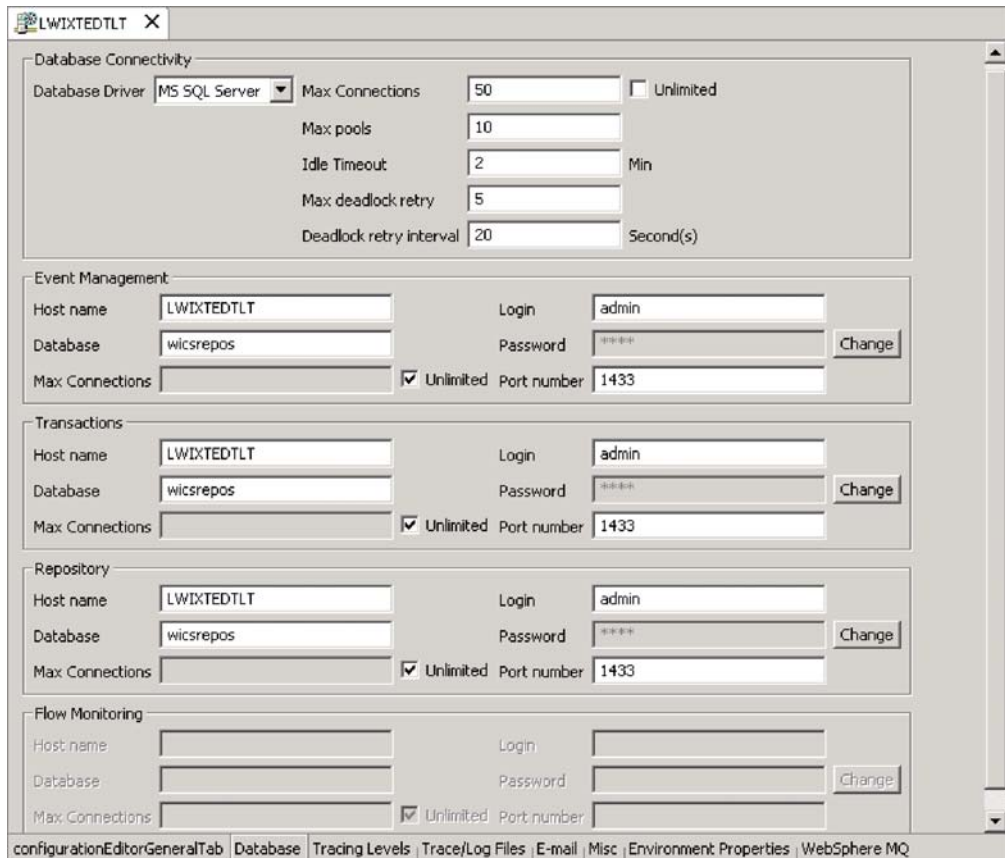


Figure 25. Database tab of Edit Configuration window

3. Change any of the database passwords:
  - a. In the section for the appropriate database (Event Management, Transactions, or Repository), click the Change button.  
A dialog box for changing the password appears.
  - b. Type the old password in the Old Password field.
  - c. Type a new password in the New Password field.  
A maximum of 30 characters is allowed.
  - d. Retype the new password in the Confirm Password field.
4. Click OK.

---

## Using System Manager

This section provides an overview of System Manager, and describes some basic administrative tasks, such as starting up, shutting down, refreshing, and setting system-wide flow control. For detailed information about using System Manager for configuration and deployment tasks, refer to the *System Implementation Guide*. This section covers the following topics:

“Starting System Manager” on page 53

“Shutting down System Manager” on page 53

“Refreshing System Manager and updating components” on page 53

“Configuring system-wide flow control” on page 54

## Starting System Manager

To start System Manager, click Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration Toolset > Administrative > System Manager. The System Manager perspective of the IBM WebSphere Studio Workbench appears (see Figure 26).

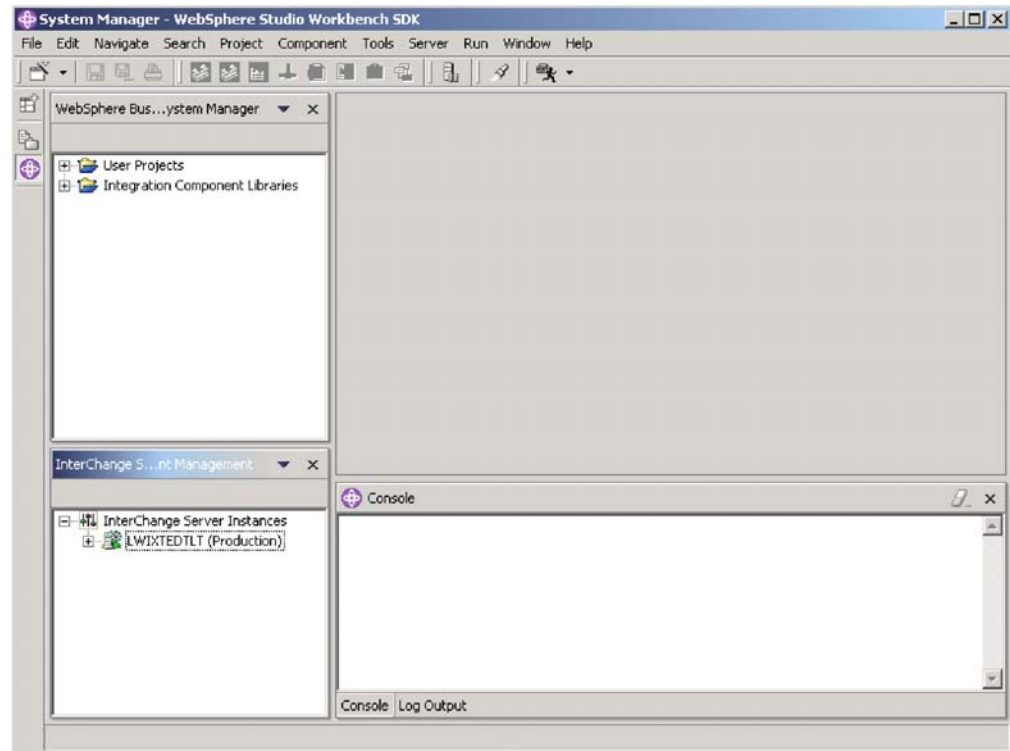


Figure 26. System Manager

## Shutting down System Manager

To shut down System Manager, select Exit from the File drop-down menu of IBM WebSphere Studio Workbench.

**Note:** Be sure to shut down any InterChange Servers before shutting down System Manager. For instructions on shutting down InterChange Servers, see “Shutting down InterChange Server” on page 49.

## Refreshing System Manager and updating components

Refreshing System Manager reloads objects from the local repository into System Manager, but does not update InterChange Server. For example, if you refresh System Manager after adding a newly created business object definition, you can add the new business object to the connector’s supported business object list and bind the connector to a collaboration port. But InterChange Server is not aware of the business object unless you reboot the server, causing the business object’s specifications to be loaded from the repository into the server’s cache.

To refresh InterChange Server, right-click the server under Server Instances, then select Refresh.

The following describes which components can be updated during system runtime:

Business objects	Not updated during runtime. The repository is read only once when InterChange Server starts up.
Collaboration object properties	Updated during runtime. For example, collaboration object trace levels take effect as soon as they are set.
Collaboration object code changes	Updated during system runtime.
Map code changes	Updated during system runtime. If mapping code is updated and recompiled, connectors must be rebound to the altered maps.

## Configuring system-wide flow control

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control for individual components, see “Configuring flow control for connectors” on page 68 or “Configuring flow control for collaboration objects” on page 74.

**Note:** Configuration changes for individual connectors or collaboration objects are dynamic, meaning they do not require ICS to be rebooted. System-wide configuration changes for flow control require ICS to be rebooted.

To monitor how flow control is working in the system, you can view the Flow Control monitor and view provided as part of System Monitor or you can view the Statistics for collaboration objects or connectors from the InterChange Server Component Management view of System Manager. For more information on using the Flow Control monitor and view in System Monitor, see “Using default monitors” on page 2 and “Using default views” on page 16. For more information on viewing the flow control from the InterChange Server Component Management view of System Manager, see “Collaboration object statistics” on page 26 or “Connector statistics” on page 27.

To configure system-wide flow control, do the following:

1. In the InterChange Server Component Management view of System Manager, right-click the ICS instance for which you want to configure flow control, then select Edit Configuration. The upper-right quadrant of System Manager changes to an editing tool in which many system properties can be changed.
2. Click the Misc tab. A dialog box appears with a Flow Control section (see Figure 27 on page 55).



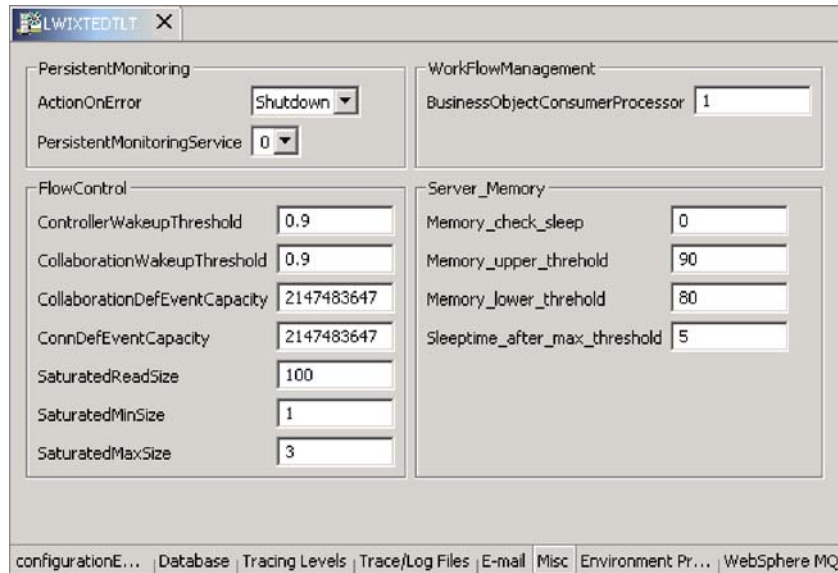


Figure 27. Edit configuration tool, Misc tab

3. In the FlowControl section, enter information in the following fields:

**ControllerWakeupThreshold:** This property applies to connector event queues. It has a decimal value ranging from 0 to 1, but not including 0 or 1. Connector event queues are always of the blocking type, meaning that if the queue is full, they do not allow new events to be added. After a queue becomes full, the connector becomes blocked. When the queue size equals or falls below the value of the connector wakeup threshold multiplied by the maximum event capacity of that connector ( $\text{CONTROLLER\_WAKEUP\_THRESHOLD} \times \text{MaxEventCapacity}$ ), the connector becomes reactivated.

**CollaborationWakeupThreshold:** This property applies to collaboration object event queues. It has a decimal value ranging from 0 to 1, but not including 0 or 1. This property applies only to blocking-type collaboration objects, meaning that it does not allow the connector to add more events to the collaboration queue. When the queue size equals or falls below the value of the collaboration object wakeup threshold multiplied by the maximum event capacity of that connector ( $\text{COLLABORATION\_WAKEUP\_THRESHOLD} \times \text{MaxEventCapacity}$ ), the connector is able to add more events to the collaboration queue for processing.

**CollaborationDefEventCapacity:** This property sets the maximum number of events you want queued for each collaboration object in the system. The range of values for this property is from 1 to 2147483647, inclusive.

**ConnDefEventCapacity:** This property sets the maximum number of events you want queued for each connector in the system. The range of values for this property is from 1 to 2147483647, inclusive.

**SaturatedReadSize:** Saturated readers attempt to process saturated events. For example, if a collaboration object queue can accept more events, the reader reads a particular number of events from the database, and then adds them to the collaboration object queue. This property reflects the maximum number of such events that can be read in one iteration of the reader.

**SaturatedMinSize:** This property applies to saturated readers, which are readers that process saturated events in the database, then add those events to the appropriate collaboration object queue. This property reflects the minimum number of threads doing these activities. The default is 1.

**SaturatedMaxSize:** This property applies to saturated readers, which are readers that process saturated events in the database, then add those events to the appropriate collaboration object queue. This property reflects the maximum number of threads doing these activities. The default is 3.

4. Save the changes you made to the ICS configuration by selecting Save *<ServerName>* from the File menu of System Manager.
5. Restart InterChange Server.

---

## Operating connectors

Operating connectors may include such tasks as starting, pausing, stopping, and shutting down connectors. For information about configuring connectors, including setting properties, supported business objects, and associated maps, see the *System Implementation Guide*.

You can start, pause, stop, and shut down connectors from either System Monitor or the InterChange Server Component Management view of System Manager.

This section covers the following topics:

“Connector states” on page 56

“Starting, stopping, and pausing connectors” on page 58

“Configuring flow control for connectors” on page 68

### Connector states

You can view the state of a connector either by logging on to System Monitor and opening a view that contains connector states or by using the InterChange Server Component Management view of System Manager. To log on to System Monitor, follow the instructions in “Logging on to System Monitor” on page 13. To use the InterChange Server Component Management view of System Manager, follow the instructions in “Connecting to an InterChange Server instance” on page 23..

The state of a connector is represented differently, depending on which tool you are using:

#### System Monitor

To see the state of connectors using System Monitor, do the following:

1. If the System Overview view is not displayed, click the System Overview link under Views in the left pane of the Web page. The System Overview Monitor appears (see Figure 11 on page 15) in the body of the Web page.

When the product is installed, the default view is set to System Overview, and the default monitor contained in that view is set to System Overview. These defaults can be changed to suit your monitoring needs. See “Using Views” on page 16 for instructions.
2. Click the triangle next to the name of the server to reveal a list of components on the system.
3. Click the triangle next to a running collaboration to reveal its associated connectors (see Figure 28 on page 57).

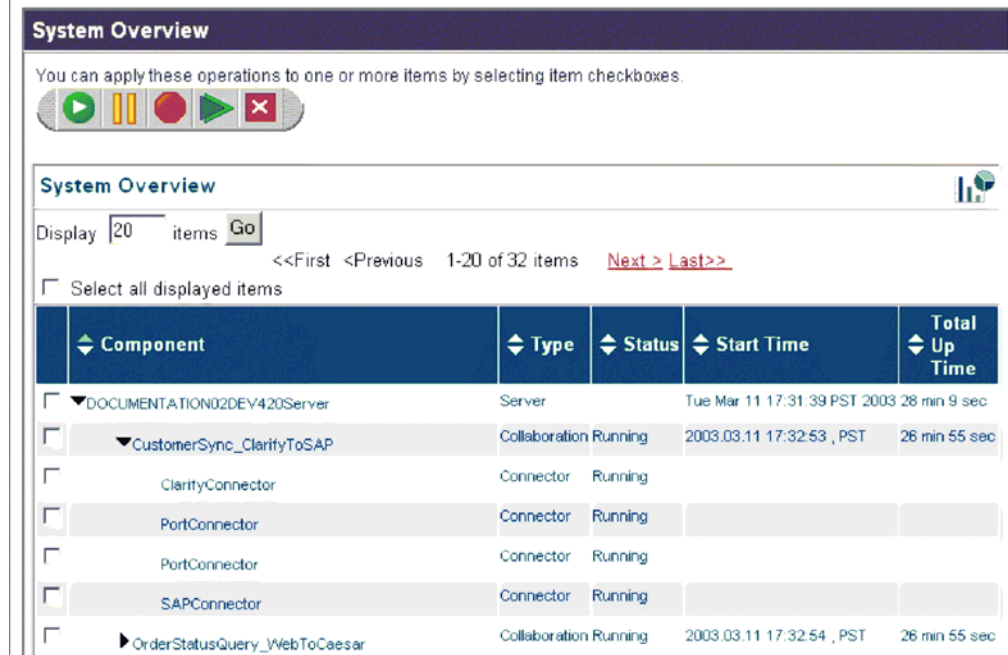


Figure 28. System Monitor, System Overview displaying connector status

### InterChange Server Component Management view of System Manager

In the InterChange Server Component Management view of System Manager, follow these steps to view the state of a connector:

1. Connect to the ICS instance that contains the connector you want to view. See "Connecting to an InterChange Server instance" on page 23 for instructions on connecting to an ICS instance.
2. Expand the ICS instance, then expand the Connectors folder.

The connectors appear under the expanded Connectors folder with different colored lights to indicate their different states.

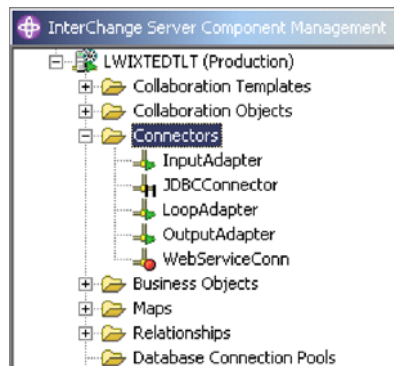


Figure 29. Connectors folder in the InterChange Server Component Management view of System Manager

Table 9 lists the connector states represented by the display color of each connector and shows what actions are being performed during that state.

Table 9. Connector States

Connector State	Subscription requests processed	Service call requests processed	Subscription deliveries processed
Active (green)	Yes	Yes	Yes
Paused (yellow)	Yes	Yes	No
In recovery or unknown (grey)			
Inactive (red)	No	No	No

## Starting, stopping, and pausing connectors

This section describe how to start, stop, and pause connectors. The following topics are covered:

“Initializing a connector” on page 58

“Running, stopping, and pausing connectors” on page 58

“Commands for changing connector states” on page 59

“Manually starting a connector” on page 59

“Shutting down a connector” on page 64

“Restarting a connector” on page 65

“Setting automatic and remote restart for a connector” on page 66

### Initializing a connector

The first time you start a connector, it must be initialized. Initializing a connector requires that you start it manually. For instructions on manually starting a connector, see “Starting a connector manually on UNIX” on page 59 or “Starting a connector manually on Windows” on page 62.

If the connector does not start, check to make sure that the command line to start it includes the current InterChange Server name. For more information on the connector’s password to InterChange Server, refer to the *System Installation Guide for UNIX* or *for Windows*.

### Running, stopping, and pausing connectors

After the connector has been initialized, you can start, stop, and pause it with System Manager or the InterChange Server Component Management view of System Manager:

#### System Monitor

1. While viewing the System Overview view (see Figure 28 on page 57), select a connector by placing check in the box to its left.
2. Select the Start, Pause, or Stop icon from the icon group in the upper-left corner of the view (see Figure 30 on page 59).



Figure 30. System Monitor, icons for starting, pausing, restarting, or shutting down components

### InterChange Server Component Management view of System Manager

From the Connectors folder in the InterChange Server Component Management view of System Manager (see Figure 29 on page 57), right-click the name of a connector, then select one of the start, pause, or stop connector options.

### Commands for changing connector states

The following list describes the commands you can use to change the connector state and describes their processing actions:

Start <i>NameConnector</i>	Starts the connector for the application <i>Name</i> if it is paused or stopped. Connectors poll the application and connector controllers read the persistent queue. Flows are processed.
Pause <i>NameConnector</i>	Pauses the connector for the application <i>Name</i> if it is running or stopped. Connectors stop polling the application and connector controllers stop reading the persistent queue. Flows are not processed.
Stop <i>NameConnector</i>	Stops the connector for the application <i>Name</i> if it is running or paused. Connectors stop polling the application and fail requests with an exception message. Connector controllers stop reading the persistent queue. Flows and requests are not processed.
Shut Down <i>NameConnector</i>	Shuts down the connector for the application <i>Name</i> . The connector's process is stopped.
Boot Up Connector Agent	Restarts the connector for the application <i>Name</i> . This action is available only if you have set the <code>OADAutoRestartAgent</code> property of the connector to True. See "Setting automatic and remote restart for a connector" on page 66.

### Manually starting a connector

The procedure for manually starting a connector depends on whether your operating system is UNIX or Windows. This section provides the following information:

- "Starting a connector manually on UNIX"
- "Starting a connector manually on Windows" on page 62

**Starting a connector manually on UNIX:** To start a connector, use the connector manager script in the *ProductDir* /bin directory with the following syntax:

```
connector_manager_connector -start
```

where *connector* is the name of the connector that you want to start. The case and spelling of this *connector* name *must* match the name of the connector's subdirectory under:

*ProductDir/connectors*

For example, the following command starts the e-Mail connector and provides the default password to InterChange Server:

```
connector_manager_EMail -start
```

The e-Mail connector has a *ProductDir/connectors* subdirectory named *EMail*. Therefore, the connector manager script must include *EMail* as the connector name (not *Email*).

The *connector\_manager\_connector* script is a wrapper for the generic connector manager script (*ProductDir/bin/connector\_manager*). This wrapper includes the following information so that you do not need to specify it:

- The name of the connector to start or stop
- Appropriate command-line options of the generic connector manager

For example:

- The SAP connector requires the *-t* command-line option. Therefore, its startup script already includes the *-t* option.
- All UNIX connectors are started with the *-b* option. Therefore, all connector startup scripts already include the *-b* option. To have a connector run in the foreground, remove the *-b* option from the generic connector manager script (*connector\_manager*).

- The name and path of the configuration file

By default, the configuration information for a connector installed for use with InterChange Server (ICS) resides in the InterChange Server repository. You can optionally use a connector configuration file that resides locally on the same machine as the connector. In some circumstances—for example, if you are using JMS—a local connector configuration file is mandatory.

For a local connector configuration file to be used, the file name and path for that file must be specified with the *-c* option as the value of the *AGENTCONFIG\_FILE* variable in the *connector\_manager\_connector* script. The wrapper passes that information when it invokes *connector\_manager*. The value can include either a literal or a relative path. If the value of the variable specifies a relative path, the startup script looks for the specified file in the directory where the product is installed.

When a connector starts up, it will first look for its configuration values in the file specified by the *AGENTCONFIG\_FILE* variable, and will then look in the InterChange Server repository for any configuration values that it did not find in the local configuration file. Values for properties in the specified local configuration file take precedence over values for the same properties in the InterChange Server repository.

If you intend to use a local configuration file for your connector, you may need to supply or change the *AGENTCONFIG\_FILE* value in the *connector\_manager\_connector* script. The necessity for this depends in part on which version of the product Installer you used to install the connector. You may have used either of two versions of the product Installer:

- A product Installer version that assumes that you are using InterChange Server. If you used this installer, it stored configuration information for the connector in the InterChange Server repository, and it also generated a local configuration file for the connector, and set the file name and path value for that configuration file in the *AGENTCONFIG\_FILE* variable in the *connector\_manager\_connector* script. The connector will look for a configuration file that has the file name and path specified in the

AGENTCONFIG\_FILE variable; if you intend to use a configuration file has a different file name and path, you must edit the AGENTCONFIG\_FILE value to match.

- A product Installer version that displays a screen requiring you to choose between InterChange Server and another “broker type.” For example, this installer version may have asked you to choose between InterChange Server and WebSphere MQ Integrator for your broker. If you chose InterChange Server, the installer stored the connector configuration in the InterChange Server repository. This installer does not create a local connector configuration file, and does not set the value of the AGENTCONFIG\_FILE variable. Therefore, if you intend to use configuration values that are stored in a local connector configuration file, you will need to create the file manually, and add its file name and path as the value of the -c option in the AGENTCONFIG\_FILE variable in the connector\_manager\_connector script.

In either of the above cases, you do not need to modify the connector\_manager\_connector script from its default installed settings if you are using the configuration stored in the InterChange Server repository.

**Important:** If you are using Installer to update an existing connector that uses a local configuration file or whose connector\_manager\_connector script has been customized, back up the original connector\_manager\_connector script and configuration file before beginning installation

To specify a local configuration file, or to change its name or path, use the ConnConfig.sh command. This command opens a graphical interface that allows you to change the connector\_manager\_connector script graphically rather than in a text editor.

The generic connector manager script calls the appropriate start\_connector.sh script, which handles the actual connector management for the connector. The IBM WebSphere Business Integration Adapter product provides a start\_connector.sh script with each connector. This start\_connector.sh script supports the options in Table 10 for starting a connector.

Table 10. Command-Line Options for the start\_connector.sh Script

Option	Additional Information
-b	This option runs the connector as a background thread; that is, no input is read from STDIN (standard input). The generic connector_manager script (called by each connector_manager_connector script) automatically specifies the -b option when it invokes the start_connector.sh script for a connector. You can remove this option from the start_connector.sh invocation to prevent a connector from being run in the background. The -b option is <i>not</i> valid on the command-line invocation of connector_manager_connector.
-fpollFrequency	<p>Poll frequency is the number of milliseconds between polling actions.</p> <ul style="list-style-type: none"> <li>• To specify the number of milliseconds, provide a value for <i>pollFrequency</i>.</li> <li>• To cause the connector to poll only when you type the value p in the connector’s Command window, specify the -fkey option.</li> <li>• If a connector is configured to processes only business object requests and not application events, polling is unnecessary; you can disable polling by specifying -fno.</li> </ul> <p>The value of this parameter overrides any repository definitions. You can specify either -fkey or -fno, but not both. The -f option is valid on the command-line invocation of connector_manager_connector. The connector manager script can pass this option to its associated start_connector.sh script.</p>

Table 10. Command-Line Options for the `start_connector.sh` Script (continued)

Option	Additional Information
<code>-tthreading_type</code>	<p>The <code>threading_type</code> parameter specifies the threading model:</p> <ul style="list-style-type: none"> <li>• <code>-tSINGLE_THREADED</code>: only a single thread accesses the application (The SAP connector uses <code>-tSINGLE_THREADED</code>.)</li> <li>• <code>-tMAIN_SINGLE_THREADED</code>: only the main thread accesses the application</li> <li>• <code>-tMULTI_THREADED</code>: multiple threads can access the application</li> </ul> <p>The <code>-t</code> option is <i>not</i> valid on the command-line invocation of <code>connector_manager_connector</code>. Specify it inside the generic <code>connector_manager</code> script, in the invocation of the <code>start_connector.sh</code> script.</p>

**Note:** The connector startup script requires the existence of the `CWSharedEnv.sh` file. If this file does not exist, the startup script generates a warning and exits. Before you attempt to run the connector startup script again, use the product Installer to create the `CWSharedEnv.sh` file. Verify that the shell startup script (such as `.cshrc`) sources the `CWSharedEnv.sh` file.

You can also use the `connector_manager_connector` script to perform the following tasks:

- Stop a connector immediately by calling the API:  
`connector_manager_connector -stop`
- Show the current status of a connector:  
`connector_manager_connector -stat`
- Terminate the process by the process Id:  
`connector_manager_connector -kill`

**Note:** The `-kill` command should not be used if Connector Agent Parallelism has been enabled, because when the `-kill` command is used the connector's dependencies will remain after the agent has been stopped. Use the `-stop` command instead.

**Note:** System Manager runs on a Windows client machine. It can monitor or stop a connector installed on a UNIX machine.

Each connector manager script has a log file with the name:  
`connector_manager_connector.log`

where `connector` is the name of the connector. Each log file contains messages generated by the `connector_manager_connector` script *and* the associated connector. For example, the `connector_manager_Oracle.log` file contains messages from the `connector_manager_Oracle` script and the IBM WebSphere ICS connector for Oracle. The log files are located in the same location as the `InterchangeSystem.log` file: `ProductDir/logs`.

**Starting a connector manually on Windows:** When you install the IBM WebSphere Business Integration Adapters on a Windows machine, a shortcut is created for each installed connector on the IBM WebSphere program menu. The connector is defined in the InterChange Server repository and is loaded when you load the repository.



Starting InterChange Server automatically initializes every connector defined in the repository. The connector is available for use whenever InterChange Server is running.

**Note:** To make a connector functional for the first time, you must configure it before you start the connector.

You can start the connector in several ways:

- click the desktop shortcut.  
Start the connector by double-clicking the desktop shortcut created as part of the installation procedure.
- Select the connector's menu option in the IBM WebSphere submenu of the Windows Start menu.
- Use a DOS Command Prompt window to execute the startup script.  
Open a DOS Command Prompt window and navigate to the appropriate connector directory. At the prompt, enter the one of the statements below, depending on whether the connector is a Java connector or a C++ connector:

**Java connector**

```
start_Sap ConnectorName InterChangeServerName
```

**C++ connector**

```
start_connector ConnectorName InterChangeServerName
```

where ConnectorName is the name of the connector and InterChangeServerName is the name of the InterChange Server instance.

**Note:** To figure out whether the connector is a Java connector or a C++ connector, navigate to *ProductDir*\documentation\wbia\_adapters\featurechecklists\versionlist.htm in your local directory, where *ProductDir* is the directory where you installed the WebSphere Business Integration Adapters product.

You can customize the startup for each connector by modifying the connector shortcut or the *start\_connector.bat* file. Use the connector startup parameters listed in Table 11 to customize the startup of a connector.

Table 11. Connector Startup Parameters

Parameter	Function
-c <i>configFile</i>	Name of the configuration file to be used during startup. If the filename specifies a relative path, the startup script looks for the file in the directory where the product is installed. This parameter is required only to use a local connector configuration file. If you are not using a local configuration file, enter the name of the configuration file used by the IBM WebSphere InterChange Server (by default, <i>InterchangeSystem.cfg</i> ).
-c	Causes the default configuration file to be used if the user-specified configuration file does not exist.
-d	Specifies the name of the C++ connector's library file, which is a dynamic link library (DLL). This DLL name does not include the .dll file extension. The startup script specifies this option for all C++ connectors.

Table 11. Connector Startup Parameters (continued)

Parameter	Function
-f <i>pollFrequency</i>	<p>Poll frequency is the number of milliseconds between polling actions.</p> <ul style="list-style-type: none"> <li>To specify the number of milliseconds, provide a value for <i>pollFrequency</i>.</li> <li>To cause the connector to poll only when you type the value p in the connector's Command Prompt window, specify the -fkey option.</li> <li>If a connector is configured to processes only business object requests and not application events, polling is unnecessary; you can disable polling by specifying -fno.</li> </ul> <p>The value of this parameter overrides any repository definitions. You can specify either -fkey or -fno, but not both.</p>
-j	Specifies that the connector is written in Java. This parameter is optional if you specify -l <i>className</i> .
-l <i>className</i>	Specifies the name of the Java connector's global class, which is an extension of the connector base class. The startup script specifies this option for all Java connectors.
-n <i>connectorName</i>	Specifies the name of the connector to start.
-p password	Specifies the password that the connector uses to access InterChange Server.
-s <i>serverName</i>	Specifies the name of the InterChange Server. This parameter is required. The name is case-sensitive.
-t	Turns on the connector property <code>SingleThreadAppCalls</code> . This property guarantees that all calls the connector framework makes to the application-specific connector code are with one event-triggered flow. The default value is false. Important: Do not change the value of this property from its shipped value. Each connector has the appropriate setting for its threading model. Specify this option only when starting a connector you created.
-x <i>connectorProps</i>	Passes application-specific connector properties to the connector. Use the format <code>prop_name=value</code> for each value you enter.

## Shutting down a connector

The generic connector manager script calls the appropriate `start_connector.sh` script, which handles the actual connector management for the connector. The IBM WebSphere ICS product provides a `start_connector.sh` script for each connector it delivers. Shutting down a connector stops the connector's processes. Before shutting down a connector, pause or stop each collaboration object that uses the connector (the collaboration must be configured to pause; see the collaboration documentation for details on how to do this). If the "Pause when critical error occurs" property has been set for a collaboration in the Collaboration General Properties window, the collaboration pauses automatically when a critical error occurs. The latest unprocessed events of such collaborations are then moved to the event submission queue.

You can perform either a "permanent" or a "temporary" shutdown of the connector. You control the type of shutdown by enabling or disabling (the default) automatic restart:

- If you have not enabled automatic restart, when you perform a shutdown action the effect is "permanent"—that is, the connector shuts down and will not restart until and unless you restart it manually at the command line or with a batch file.

- If you have enabled automatic restart, the shutdown action is temporary, and you can restart the connector by using the Boot Up Connector Agent action in System View.

For instructions on enabling or disabling automatic restart, see “Setting automatic and remote restart for a connector” on page 66.

To shut down a connector:

#### UNIX

Use the following command: `connector_manager_<connector_name> -stop`.

#### Windows

Use System Monitor to shut down a connector.

Instructions for shutting down a connector using System Monitor depend on which version of System Monitor you are using:

#### System Monitor

1. From the System Overview view, select the collaboration object of the connector you want to shut down by placing a check in the box to its left, then click the Pause icon from the upper-left corner of the view (see Figure 30 on page 59). Do this for each collaboration associated with the connector.
2. Select the connector you want to shut down by placing check in the box to its left, then click the Shutdown icon from the upper-left corner of the view (see Figure 30 on page 59).

#### InterChange Server Component Management view of System Manager

1. From the expanded Collaboration Objects folder in the InterChange Server Component Management view of System Manager, (see Figure 33 on page 71), right-click the collaboration object associated with the connector, then choose Pause. Do this for each collaboration associated with the connector.  
The color on the collaboration object icon turns to yellow.
2. From the expanded Connectors folder in the InterChange Server Component Management view of System Manager, right-click the connector, then select Shut Down.

**Attention:** Do not use the Ctrl+C key sequence to shut down a connector. Doing so prevents the connector from shutting down in an orderly manner. In addition, if you do use the Ctrl+C key sequence, or if you use “q,” or other manual methods to perform the shutdown, and OAD is enabled, OAD will immediately restart the connector.

#### Restarting a connector

This action is used to restart the connector after you have used the Shut Down Connector action in either System Monitor or the InterChange Server Component Management view of System Manager. This action is available only if you have enabled automatic and remote restart for the connector (see “Setting automatic and remote restart for a connector” on page 66).

Instructions for restarting a connector depend on which tool you are using:

## System Monitor

1. From the System Overview view (see Figure 28 on page 57), place a check in the box to the left of the connector you want to restart.
2. Click the Restart Agent icon from the upper-left corner of the view (see Figure 30 on page 59).

**InterChange Server Component Management view of System Manager** From the expanded Connectors folder in the InterChange System Component Management view of System Manager, (see Figure 29 on page 57), right-click the connector you want to restart, then click Boot <name\_of\_connector>.

## Setting automatic and remote restart for a connector

With the WebSphere MQ-triggered Object Activation Daemon (OAD), you can enable a connector to support the automatic-and-remote-restart feature, which allows the connector to handle the following conditions:

- Availability: restart automatically after the connector has been shut down
- Serviceability: start or restart a remote connector agent from System Manager

### High Availability

If a connector is already a member of an HA group, the automatic restart property would be redundant and should be disabled

To set up automatic and remote restart for a connector, take the following steps:

- Ensure that WebSphere MQ is installed
- Start the MQ Trigger Monitor
- Configure an adapter for the automatic and remote restart

**Installing the OAD software:** Use of the MQ-triggered OAD requires installation of the MQ-trigger Monitor and the configuration of certain queues. This monitor is installed as part of the WebSphere MQ software. These queues are created and configured by a special `mqtriggersetup` script. These steps are described in the *System Installation Guide for UNIX or for Windows*.

**Important:** The WebSphere MQ triggering monitor must exist on the machine on which the connector agent resides. If multiple connector agents reside on a single machine, only one MQ Trigger Monitor needs to exist.

**Starting the MQ-triggered OAD:** To start the MQ-triggered OAD, you must start MQ Trigger Monitor, which can be done in either of the following ways:

- Explicitly start MQ Trigger Monitor with the appropriate startup script
- Install MQ Trigger Monitor as a Windows service

For information on how to start MQ Trigger Monitor in either of these two ways, refer to the "Starting MQ-triggered OAD" section in the *Installation Guide for UNIX or for Windows*.

**Enabling connectors for MQ-triggered OAD:** To enable automatic and remote restart for a connector, you set OAD connector configuration properties as shown in Table 12 on page 67.

Table 12. Configuring standard properties in Connector Configurator

Name	Possible Values	Description	Default Values
OADAutoRestartAgent	true or false	If set to true, the MQ-triggered OAD will automatically attempt to restart the connector after an abnormal shutdown. It can also be used to start the connector agent remotely. This value is dynamic.	false
OADMaxNumRetry	Number	Number of maximum tries that activation will retry	10,000
OADRetryTimeInterval	Minutes	Number of minutes between each retry. If the connector agent does not start in this time interval, another attempt to restart the agent will be made.	10

To set the OAD properties, you must start Connector Configurator for the connector, using the following steps:

1. Start ICS.
2. From System Manager, double-click the connector. This opens Connector Configurator.
3. From the Standard Properties tab, set the standard properties shown in Table 12.

From within Connector Configurator, you can take any of the following actions:

- Initialize the connector for MQ-triggered OAD
- Toggle the MQ-triggered OAD for a connector

**Initializing a connector for MQ-triggered OAD:** The following steps describe how to enable automatic and remote restart for a connector for the first time.

1. Set the OADAutoRestartAgent property to True.
2. Set any of the other desired OAD properties in Table 12.
3. Save the OAD properties in Connector Configurator.

**Toggling automatic and remote restart:** Changing the value of the OADAutoRestartAgent property from True to False toggles the automatic-and-remote-restart feature on and off. This connector property is dynamic; that is, you do not need to restart ICS for the change to take effect. Therefore, when you set OADAutoRestartAgent to False, automatic and remote restart is disabled. When you set this property to True, automatic restart is enabled.

If you shut down the connector agent while the automatic-and-remote-restart feature is enabled, you perform a temporary shutdown. The response of the connector depends on the method you use to shutdown the connector, as follows:

- If you shut down the connector from its connector-startup window (by typing "q" or Ctrl+C), the connector agent shuts down and MQ-triggered OAD automatically restarts the connector.

- If you shut down the connector from within System Manager (by clicking the Shutdown Agent button), the connector agent shuts down. However, the MQ-triggered OAD is not able to automatically restart the connector. You must restart the agent from within System Manager (by clicking the Reboot Agent button).

However, if you disable the automatic-and-remote-restart feature and then shut down the connector agent, you perform a permanent shutdown; that is, you must manually restart the connector.

## Configuring flow control for connectors

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control system-wide, see “Configuring system-wide flow control” on page 54. This section describes how to configure flow control for connectors.

**Note:** Configuration changes for individual connectors or collaboration objects are dynamic, meaning they do not require ICS to be rebooted. System-wide configuration changes for flow control require ICS to be rebooted.

To monitor how flow control is working in the system, you can view the Flow Control monitor and view provided as part of System Monitor or you can view the Statistics for collaboration objects or connectors from the InterChange Server Component Management view of System Manager. For more information on using the Flow Control monitor and view in System Monitor, see “Using default monitors” on page 2 and “Using default views” on page 16. For more information on viewing the flow control from the InterChange Server Component Management view of System Manager, see “Collaboration object statistics” on page 26 or “Connector statistics” on page 27.

To configure flow control for a connector, do the following:

1. In System Manger, navigate to the connector for which you want to configure flow control, then double-click that connector. Connector Configurator opens (see Figure 31 on page 69).

	Property	Value	Update Method
1	AgentConnections	1	component restart
2	AgentTraceLevel	0	dynamic
3	ApplicationName	ClarifyConnector	component restart
4	BrokerType	ICS	connector restart
5	CharacterEncoding	ascii7	component restart
6	ConcurrentEventTriggeredFlows	1	component restart
7	ControllerStoreAndForwardMode	true	dynamic
8	ControllerTraceLevel	0	dynamic
9	DeliveryTransport	MQ	component restart
10	JvmMaxHeapSize	128m	component restart
11	JvmMaxNativeStackSize	128k	component restart
12	JvmMinHeapSize	1m	component restart
13	Locale	en_US	component restart
14	LogAllInterchangeEnd	false	component restart
15	MaxEventCapacity	2147483647	dynamic
16	MessageFileName	ClarifyConnector.txt	component restart
17	OADAutoRestartAgent	false	dynamic
18	OADMaxNumRetry	1000	dynamic
19	OADRetryTimeInterval	10	dynamic
20	PollEndTime	HH:MM	component restart
21	PollFrequency	10000	dynamic
22	PollStartTime	HH:MM	component restart
23	RepositoryDirectory	<REMOTE>	Agent restart
24	RestartRetryCount	3	dynamic
25	RestartRetryInterval	1	dynamic
26	WireFormat	CwBO	agent restart

Figure 31. Connector Configurator, Standard Properties tab

2. In the Standard Properties tab, click in the Value cell of the MaxEventCapacity property.
3. Change the value to represent the maximum number of events you want queued for a connector. The valid range of values for this property is from 1 to 2147483647.
4. Click Save > to Project from the File drop-down menu. The following message appears in the bottom section of Connector Configurator: Connector '<name\_of\_connector>' is saved successfully.

## Operating collaboration objects

Operating collaboration objects may include such tasks as running, pausing, stopping, and shutting down collaboration objects. For information about configuring collaboration objects, see the *System Implementation Guide*.

You can run, pause, stop, and shut down collaboration objects from either System Monitor or the InterChange Server Component Management view of System Manager.

This section covers the following topics:

“Collaboration object states” on page 69

“Starting, stopping, and pausing collaboration objects” on page 72

“Configuring collaboration object run-time properties” on page 72

### Collaboration object states

You can view the state of a collaboration object either by logging on to System Monitor and opening a view that contains collaboration object information or by using the InterChange Server Component Management view of System Manager. To log on to System Monitor, follow the instructions in “Logging on to System

Monitor” on page 13.. To use the InterChange Server Component Management view of System Manager, follow the instructions in “Connecting to an InterChange Server instance” on page 23..

The state of a collaboration object is represented differently, depending on which tool you are using:

### System Monitor

To see the state of collaboration objects using System Monitor, do the following:

1. If the System Overview view is not displayed, click the System Overview link under Views in the left pane of the Web page. The System Overview Monitor appears (see Figure 11 on page 15) in the body of the Web page.

When the product is installed, the default view is set to System Overview, and the default monitor contained in that view is set to System Overview. These defaults can be changed to suit your monitoring needs. See “Using Views” on page 16 for instructions.

2. Click the triangle next to the name of the server to reveal a list of components on the system. All collaboration objects are listed along with their status, start time, and total up time (see Figure 32).

The screenshot shows the 'System Overview' interface. At the top, there is a header 'System Overview' and a toolbar with icons for play, pause, stop, refresh, and delete. Below the toolbar, there is a search bar with 'Display 20 items' and a 'Go' button. A navigation bar shows '<<First <Previous 1-20 of 32 items Next > Last>>'. A checkbox 'Select all displayed items' is present. The main content is a table with the following columns: Component, Type, Status, Start Time, and Total Up Time.

Component	Type	Status	Start Time	Total Up Time
▼ DOCUMENTATION02DEV420Server	Server		Tue Mar 11 17:31:39 PST 2003	28 min 9 sec
▼ CustomerSync_ClarifyToSAP	Collaboration	Running	2003.03.11 17:32:53 , PST	26 min 55 sec
ClarityConnector	Connector	Running		
PortConnector	Connector	Running		
PortConnector	Connector	Running		
SAPConnector	Connector	Running		
▶ OrderStatusQuery_WebToCaesar	Collaboration	Running	2003.03.11 17:32:54 , PST	26 min 55 sec

Figure 32. System Monitor, System Overview displaying collaboration object status

### InterChange Server Component Management view of System Manager

In the InterChange Server Component Management view of System Manager, follow these steps to view the state of a collaboration object:

1. Connect to the ICS instance that contains the collaboration object you want to view. See “Connecting to an InterChange Server instance” on page 23 for instructions on connecting to an ICS instance.
2. Expand the ICS instance, then expand the Collaboration Objects folder.



The collaboration objects appear under the expanded Collaboration Objects folder with different colored lights to indicate their different states (see Figure 33).

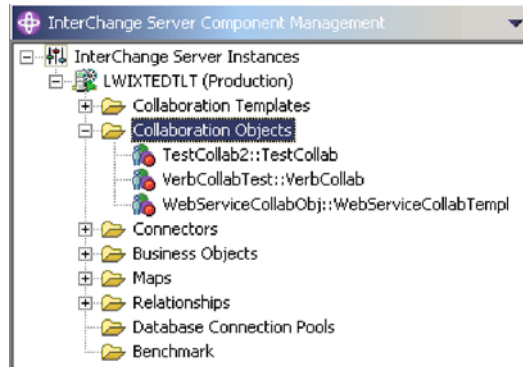


Figure 33. Collaboration Objects folder in the InterChange Server Component Management view of System Manager

Table 13 describes the collaboration object states which are viewable from System Monitor and the Collaboration Objects folder in the InterChange Server Component Management view of System Manager:

Table 13. Collaboration object states

Collaboration object state	Description
Start	Starting a collaboration object causes it to subscribe to its triggering business objects and to process them as they arrive. If you stop and then restart InterChange Server, collaboration objects in the Start state automatically start running when InterChange Server comes back up.
Pause	Pausing a collaboration prevents it from receiving new flow initiators. The collaboration completes all of the current processing, then enters an idle state.  A connector maintains its subscription information; therefore, it continues to send flow initiators to the connector queues. The collaboration processes these when it is resumed.  To resume collaboration execution, click Start in System Monitor or in the Collaboration Object menu of the InterChange Server Component Management view of System Manager.
Stop	Stopping a collaboration causes it to unsubscribe to business objects. The collaboration completes all of the current processing, then becomes inactive. Unlike the Pause command, the Stop command causes connectors to stop sending business objects to the collaboration.  To properly stop a collaboration without losing any flows, first stop the associated connectors from polling, allow all flows to process, then stop the collaboration.

Table 13. Collaboration object states (continued)

Collaboration object state	Description
Shut Down	Shutting down a collaboration immediately ends processing of current flows. When the collaboration is restarted, the system recovers by processing those flows that were interrupted by the shut down and recovering those flows waiting in the queue. This recovery is not immediate, so prepare to wait while the system completes the recovery interval.

**Note:** When you stop or shut down a collaboration object that is part of a collaboration group, all collaborations in the group stop or shut down. If any member of a collaboration group fails to start up or has a state change failure, the collaboration group is rolled back to the initial state (deactivated or stopped).

## Starting, stopping, and pausing collaboration objects

To make a collaboration object functional for the first time, you must first configure it then start it. See “Configuring collaboration object run-time properties” on page 72 for more information on configuring collaborations. Depending on which tool you are using, you run, stop, and pause collaboration objects in the following way:

### System Monitor

1. While viewing the System Overview view (see Figure 28 on page 57), select a collaboration object by placing check in the box to its left.
2. Select the Start, Pause, or Stop icon from the icon group in the upper-left corner of the view (see Figure 30 on page 59).

### InterChange Server Component Management view of System Manager

From the Collaboration Objects folder of the InterChange Server Component Management view, right-click a collaboration object, then select start, pause, or stop.

## Configuring collaboration object run-time properties

This section describes some aspects of collaboration object behavior that are configurable in a production environment and contains the following topics:

“Setting collaboration object general properties” on page 73

“Configuring collaboration objects to process concurrent event-triggered flows” on page 74

“Configuring flow control for collaboration objects” on page 74

For information about the following tasks, see the *System Implementation Guide*:

- Creating a Collaboration Object
- Configuring Collaboration-Specific Properties
- Binding the Ports of the Collaboration
- Setting the Effective Transaction Level and Other General Properties

## Setting collaboration object general properties

To open the Collaboration Properties window and change values for general properties of a collaboration object, do this:

1. From the expanded Collaboration Objects folder in the InterChange Server Component Management view of System Manager, right-click a collaboration object and select Properties.
2. In the Properties dialog box, choose the Collaboration General Properties tab. The following dialog displays:

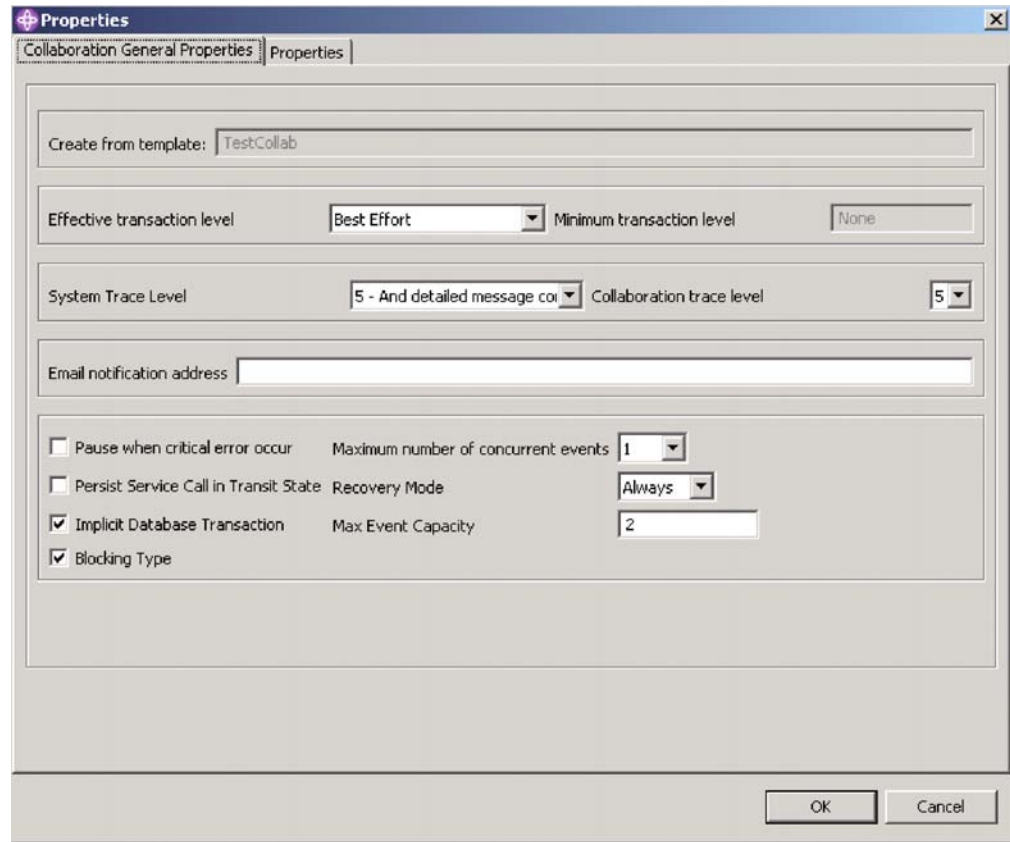


Figure 34. Properties dialog box, Collaboration General Properties tab

The dialog box shows the template from which the collaboration object was generated and the minimum transaction level that was specified in the collaboration template.

The dialog enables you to make settings for the following:

- Effective Transaction Level  
See "Setting the Effective Transaction Level and Other General Properties" in the *System Implementation Guide*.
- System Trace Level  
See "Configuring logging and tracing" on page 129.
- Collaboration Trace Level  
See "Configuring the collaboration object trace level" on page 135.
- E-mail Notification Address

See “Configuring e-mail notification at the collaboration object level” on page 125.

- Pause When Critical Error Occurs  
See “Critical errors” on page 148.
- Persist Service Call In Transit State  
See Chapter 3, “Troubleshooting the system,” on page 117, and see the *Collaboration Development Guide*.
- Implicit Database Transaction  
See “Configuring Transaction Bracketing” in the *System Implementation Guide*.
- Blocking Type  
See “Blocking type” in the *System Implementation Guide*.
- Maximum Number of Concurrent Events  
See “Configuring collaboration objects to process concurrent event-triggered flows” on page 74.
- Recovery Mode  
See “Deferred recovery of collaboration events” on page 147.
- Max Event Capacity  
See “Configuring flow control for collaboration objects.”

### **Configuring collaboration objects to process concurrent event-triggered flows**

For detailed information about processing concurrent events, see the *System Implementation Guide*.

**Tip:** Processing concurrently triggered events in collaborations requires additional system resources. To maximize performance, ensure that system resources used to handle concurrent events are not idle. For example, do not set the value for the maximum concurrent triggered-event processing option to 10 if the collaboration queue is set to process a maximum of four events.

To set the maximum number of concurrent flows for a collaboration:

1. From the expanded Collaboration Objects folder in the InterChange Server Component Management view, right-click the collaboration object that you want to change, then select Properties. The Properties dialog box appears (see Figure 34 on page 73).
2. In the Collaboration General Properties tab, enter a value in the “Maximum number of concurrent events” field.
3. Click OK to save your changes and close the window.
4. Restart the collaboration for changes to take effect.

### **Configuring flow control for collaboration objects**

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control system-wide, see “Configuring system-wide flow control” on page 54. This section describes how to configure flow control for collaboration objects.

**Note:** Configuration changes for individual connectors or collaboration objects are dynamic, meaning they do not require ICS to be rebooted. System-wide configuration changes for flow control require ICS to be rebooted.

To monitor how flow control is working in the system, you can view the Flow Control monitor and view provided as part of System Monitor or you can view the Statistics for collaboration objects or connectors from the InterChange Server Component Management view of System Manager. For more information on using the Flow Control monitor and view in System Monitor, see “Using default monitors” on page 2 and “Using default views” on page 16. For more information on viewing the flow control from the InterChange Server Component Management view of System Manager, see “Collaboration object statistics” on page 26 or “Connector statistics” on page 27.

To configure flow control for a collaboration object, do the following:

1. From the expanded Collaboration Objects folder in the InterChange Server Component Management view of System Manager, right-click the collaboration object for which you want to create flow control, then select Properties from the drop-down menu. The Properties dialog box appears (see Figure 34 on page 73).
2. In the Collaboration General Properties tab, edit the value in the Max Event Capacity field to represent the maximum number of events you want queued for a collaboration object. The valid range of values for this property is from 1 to 2147483647.
3. Click OK. The property is changed immediately.

### **Reconfiguring the timeout attribute for long-lived business processing**

Long-lived business processing enables collaboration objects to be deployed as a long-lived business processes. If a collaboration object has been configured with long-lived business processing, the service call timeout values can be reconfigured during runtime. For more information about developing a collaboration object with long-lived business processing, see the *Collaboration Development Guide*.

To reconfigure the service call timeout values of a collaboration with long-lived business processing, do the following:

1. From the expanded Collaboration Objects folder in the InterChange Server Component Management view, right-click the collaboration object whose service call timeout value you want to edit, then click Properties. The Properties dialog box appears.
2. From the Properties tab, locate the property that represents the service call timeout value you want to change, then click in the value field. When the property becomes highlighted, the value can be edited.

**Note:** The name of the service call timeout configuration property may be something like, “CreateTimeout” or “RetrieveTimeout,” but since there is no naming convention for this property, you may have to contact the person who developed the collaboration, if the name of the service call timeout configuration property is not immediately apparent.

3. Edit the value so that it represent the number of timeout minutes allowed.

**Note:** The Value field must contain an integer greater than 0. If it contains a 0 or is left blank, the waittime is equal to infinity. If it contains non-numerical values, it will trigger a collaboration runtime exception.

4. Click OK. Your changes take place immediately, without the need to restart InterChange Server.

---

## Operating maps

You can start and stop maps from either System Monitor or the InterChange Server Component Management view of System Manager.

This section covers the following topics:

“Map states”

“Starting and stopping maps” on page 77

### Map states

You can view the state of a map either by logging on to System Monitor and opening a view that contains map status or by using the InterChange Server Component Management view of System Manager. To log on to System Monitor, follow the instructions in “Logging on to System Monitor” on page 13. To use the InterChange Server Component Management view of System Manager, follow the instructions in “Connecting to an InterChange Server instance” on page 23..

The state of a map is represented differently, depending on which tool you are using:

#### System Monitor

To see the state of map using System Monitor, do the following: If the System Overview view is displayed, click the Maps and Relationships link under Views in the left pane of the Web page. The Map Status and Relationship Status monitors appear (see Figure 11 on page 15) in the body of the Web page. When the product is installed, the default view is set to System Overview, and the default monitor contained in that view is set to System Overview. These defaults can be changed to suit your monitoring needs. See “Using Views” on page 16 for instructions.

#### InterChange Server Component Management view of System Manager

In the InterChange Server Component Management view of System Manager, follow these steps to view the state of a map:

1. Connect to the ICS instance that contains the map you want to view. See “Connecting to an InterChange Server instance” on page 23 for instructions on connecting to an ICS instance.
2. Expand the ICS instance, then expand the Maps folder. The maps appear under the expanded Maps folder with different colored lights to indicate their different states (see Figure 35 on page 77).

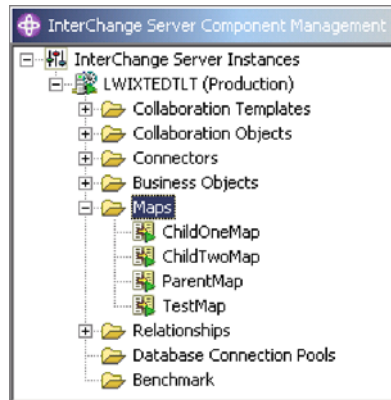


Figure 35. Maps folder in the InterChange Server Component Management view of System Manager

Table 14 lists the map states represented by the display color on the map icon and describes what actions can be performed during that state. For more information about maps, refer to the *Map Development Guide*.

Table 14. Map States

Map State/Traffic Light	Description
Active (green)	Map is ready to run and available for use in the IBM WebSphere ICS system.
Inactive (red)	Map is not ready to run or available for use in the IBM WebSphere ICS system.
Unknown (do not display)	Map has not been compiled or the map class is missing. The map is not ready to run or available for use in the IBM WebSphere ICS system. Compiling the map in Map Designer puts the map in active state, while saving the map puts the map in unknown state. Maps with unknown state display in the object browser map tree.

## Starting and stopping maps

Maps define the transfer (or transformation) of data between the source and destination business objects. In the IBM WebSphere ICS environment, data is mapped from an application-specific business object to a generic business object or from a generic business object to an application-specific business object. For detailed information about how maps are used in the IBM WebSphere ICS system, refer to the *Map Development Guide*.

This section describes the how to start and stop maps. For information about additional tasks in using maps, including map compilation, map properties, data validation levels, explicit and implicit transaction bracketing, and map instance reuse, see the *System Implementation Guide*.

### Map activation

For a map to be executable, it must first be activated. Map Designer automatically starts a map when it successfully compiles the map. However, other changes to the map might require that you explicitly stop and restart the map for the change to take effect. See Appendix B, “Requirements for restarting IBM WebSphere Business Integration system components,” on page 179 to find out what changes require maps to be stopped and restarted.

## Starting and stopping maps

Depending on which tool you are using, you start and stop maps in the following way:

### System Monitor

1. If the System Overview view is displayed, click the Maps and Relationships link under Views in the left pane of the Web page. The Map Status and Relationship Status monitors appear (see Figure 11 on page 15) in the body of the Web page.
2. Select the Start or Stop icon from the icon group in the upper-left corner of the view (see Figure 30 on page 59).

### InterChange Server Component Management view of System Manager

1. From the expanded Maps folder in the InterChange Server Component Management view (see Figure 35 on page 77), right-click a map, then select either the Start <name\_of\_map> or Stop <name\_of\_map> option.

---

## Operating relationships

You can start and stop relationships from either System Monitor or the InterChange Server Component Management view of System Manager.

This section covers the following topics:

“Relationship states” on page 78

“Starting and stopping relationships” on page 79

### Relationship states

You can view the state of a relationship either by logging on to System Monitor and opening a view that contains relationship status or by using the InterChange Server Component Management view of System Manager. To log on to System Monitor, follow the instructions in “Logging on to System Monitor” on page 13. To use the InterChange Server Component Management view of System Manager, follow the instructions in “Connecting to an InterChange Server instance” on page 23..

The state of a relationship is represented differently, depending on which tool you are using:

#### System Monitor

To see the state of relationship using System Monitor, do the following: If the System Overview view is displayed, click the Maps and Relationships link under Views in the left pane of the Web page. The Map Status and Relationship Status monitors appear (see Figure 11 on page 15) in the body of the Web page. When the product is installed, the default view is set to System Overview, and the default monitor contained in that view is set to System Overview. These defaults can be changed to suit your monitoring needs. See “Using Views” on page 16 for instructions.

#### InterChange Server Component Management view of System Manager



In the InterChange Server Component Management view of System Manager, follow these steps to view the state of a relationship:

1. Connect to the ICS instance that contains the relationship you want to view. See “Connecting to an InterChange Server instance” on page 23 for instructions on connecting to an ICS instance.
2. Expand the ICS instance, then expand the Relationships folder, then expand either the Dynamic or Static folder. The relationships appear under either the expanded Dynamic folder or the expanded Static folder, and have different colored lights to indicate their different states (see Figure 36).

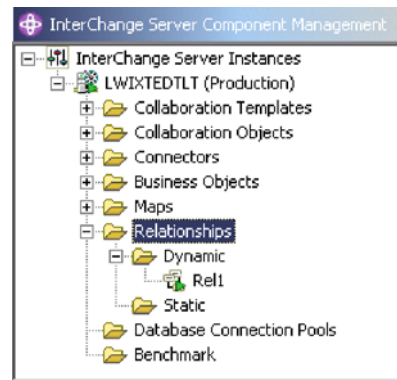


Figure 36. Relationships folder in the InterChange Server Component Management view of System Manager

Table 15 lists the relationship states represented by the display color and describes what actions can be performed during that state.

Table 15. Relationship States

Relationship state/Display color	Description
Active (green)	Relationship is ready to run and available for use in the IBM WebSphere ICS system. To use Relationship Manager on a relationship, the relationship must be in the active state.
Inactive (red)	Relationship is not ready to run or available for use in the IBM WebSphere ICS system. This state is entered when the relationship is stopped, where all current jobs in queue are completed and no new jobs are accepted. To modify a relationship definition, it must be in this state.
Unknown (grey)	Relationship does not have a compatible runtime schema. To create a compatible runtime schema, from the Relationship Designer, save the relationship with the Create runtime schema option selected. The state changes to Inactive, at which point the relationship can then be started.

## Starting and stopping relationships

Relationships are used to establish associations between business object attributes that cannot easily be mapped. The tool used for creating relationships is Relationship Designer. For more information about Relationship Designer, see the *Map Development Guide*.

When you expand the Relationships folder in the InterChange Server Component Management view of System Manager, two subfolders appear: Dynamic and Static.

- Dynamic relationship—a relationship whose runtime data changes frequently; that is, its relationship tables have frequent Insert, Update, or Delete operations. All relationships are dynamic by default.
- Static relationship—a relationship whose runtime data undergoes very minimal change; that is, its relationship tables have very few Insert, Update, or Delete operations. For example, because lookup tables store information such as codes and status values, their data very often is static. Such tables make good candidates for being cached in memory.

This section describes the following topics:

“Relationship activation” on page 80

“Starting and stopping relationships”

“Relationship table caching”

### Relationship activation

For a relationship to be executable, it must be activated. However, you cannot modify a relationship when it is active. Therefore, you must stop the relationship, make the change to the relationship, and then restart the relationship. See Appendix B, “Requirements for restarting IBM WebSphere Business Integration system components,” on page 179 to find out what changes require relationships to be restarted.

### Starting and stopping relationships

Depending on which tool you are using, you start and stop relationships in the following way:

#### System Monitor

1. From the System View window (see Figure 11 on page 15), select Show Maps and Relationships from the View drop-down menu. The Maps and Relationships columns appear next to Collaborations and Connectors.
2. Right-click a relationship, then select either the Start <name\_of\_relationship> or Stop <name\_of\_relationship> option.

#### InterChange Server Component Management view of System Manager

From the expanded Dynamic or Static folder in the InterChange Server Component Management view of System Manager (see Figure 36 on page 79), right-click the name of a relationship, then select Start <relationship\_name> or Stop <relationship\_name>.

### Relationship table caching

As part of the design process of a static relationship, a developer can indicate whether the relationship’s tables are to be cached in memory. A static relationship is one whose data does not change frequently. If the developer has indicated that the static relationship’s tables can be cached, you can control whether to enable caching from System Manager. System Manager lists all static relationships in the folder labelled Static under the Relationships folder.

**Note:** For information on how to design a static relationship so that its tables to be cached in memory, see the *Map Development Guide*.

**Enabling Caching:** To enable relationship table caching for a static relationship:

1. Expand the Relationships folder in System Manager.
2. Expand the Static folder in the object browser to locate the static relationship whose tables you want to be cached.
3. Right-click the static relationship to determine its current cached state. If the Cached option appears with no check mark to the left, caching for that relationship is currently disabled. Choose Cached from the context menu to enable caching.

When the Cached option appears with a check mark to the left, InterChange Server reads the relationship tables into memory the next time the runtime data is accessed.

**Disabling Caching:** To disable relationship table caching for a static relationship:

1. Expand the Relationships folder in System Manager.
2. Expand the Static folder in the object browser to locate the static relationship whose tables you do not want to be cached.
3. Right-click the static relationship to determine its current cached state. If the Cached option appears with a check mark to the left, caching for that relationship is currently enabled. Choose Cached from the context menu to disable caching.

When the Cached option appears with no check mark to the left, InterChange Server reads runtime data from the tables in the relationship database.

**Reloading the Cached Tables:** You can tell InterChange Server to reread the relationship's tables into memory with the Reload feature, as follows:

1. Expand the Relationships folder in System Manager.
2. Expand the Static folder in the object browser to locate the static relationship whose tables you want to be reloaded.
3. Right-click the static relationship to determine its current cached state. If the Cached option appears with a check mark to the left, caching for that relationship is currently enabled. Therefore, the Reload option is enabled.
4. Choose Reload from the context menu to reload the static relationship's tables.

When you choose this option, InterChange Server reloads the cached relationship tables by rereading the tables from the relationship database into memory. This option is useful when the static relationship's tables are updated directly in the database through SQL statements. To get the more current version of the tables into cache, choose the Reload option.

**Tracing Cached Tables:** To tell InterChange Server to log a trace message each time it loads and unloads relationship tables in memory, set the RELATIONSHIP.CACHING configuration parameter to five (5) in the TRACING section of the InterchangeSystem.cfg file:

```
RELATIONSHIP.CACHING=5
```

ICS routes these messages to the trace file (if one is configured). By default, ICS does *not* generate trace messages when it loads and unloads the relationship tables. Trace levels less than five (0-4) do not produce messages either.

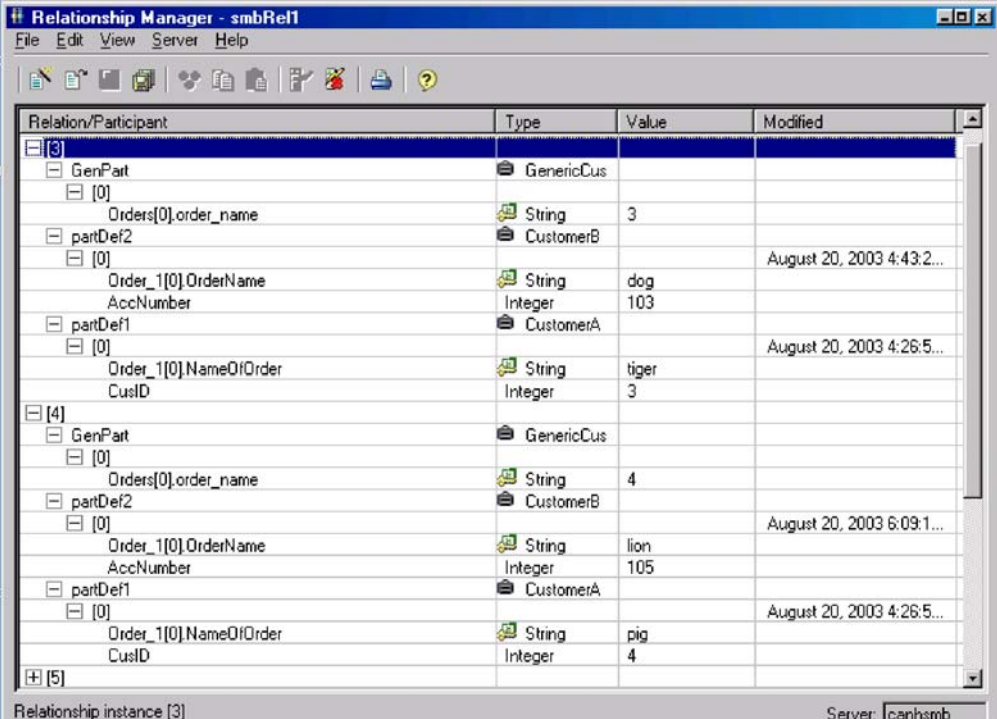
## Using Relationship Manager

Relationship Manager allows you to view and perform operations on relationship runtime data, including participants and their data. For background information about relationships, see the *Map Development Guide*.

You create relationship definitions with Relationship Designer. At runtime, instances of the relationships are populated with the data that associates information from different applications. This relationship instance data is created when the maps that use the relationships execute. The data is stored in the relationship tables specified in the relationship definition. Relationship Manager provides a graphical interface to interact with the relationship tables regardless of the database vendor.

For each relationship instance, Relationship Manager displays a hierarchical listing of its participant definitions and participant instances, which are a set of key and non-key attributes. The relationship tree also provides detailed information about each of the participants in the relationship instance such as the type of entity, its value, and the date it was last modified. A relationship instance ID is automatically generated when the relationship instance is saved in the relationship table. Relationship Manager displays this instance ID at the top level of the relationship tree.

Figure 37 shows a sample of a relationship tree in Relationship Manager for an identity relationship.



Relation/Participant	Type	Value	Modified
[3]			
GenPart	GenericCus		
[0]			
Orders[0].order_name	String	3	
partDef2	CustomerB		
[0]			August 20, 2003 4:43:2...
Order_1[0].OrderName	String	dog	
AccNumber	Integer	103	
partDef1	CustomerA		
[0]			August 20, 2003 4:26:5...
Order_1[0].NameOfOrder	String	tiger	
CusID	Integer	3	
[4]			
GenPart	GenericCus		
[0]			
Orders[0].order_name	String	4	
partDef2	CustomerB		
[0]			August 20, 2003 6:09:1...
Order_1[0].OrderName	String	lion	
AccNumber	Integer	105	
partDef1	CustomerA		
[0]			August 20, 2003 4:26:5...
Order_1[0].NameOfOrder	String	pig	
CusID	Integer	4	
[5]			

Relationship instance [3] Server: canhsmb

Figure 37. Relationship Manager, relationship tree

You can use Relationship Manager to work on entities at all levels: the relationship instance, participant instance, and attribute levels. For example, you can use Relationship Manager to:

- Create and delete relationship instances

- Modify the contents of a relationship instance, such as adding and deleting participants
- Add and save a participant's data, load a participant's data from or save data to a file, and copy and paste the data of a participant from another relationship into the relationship instance, thus creating a new participant (as long as the participant types are identical)
- Activate and deactivate participants
- Retrieve participants based on instance IDs, business object attribute values, or data
- Filter a participant's activity within a time interval
- Salvage a situation when problems with data arise. For example, when corrupt or inconsistent data from a source application has been sent to the generic and destination application relationship tables, you can use Relationship Manager to rollback (or clean up) the data back to a point of time when you know the data is reliable.

## Starting Relationship Manager

Do one of the following to start Relationship Manager:

- Select **Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration Toolset > Administrative > Relationship Manager**
- In Relationship Designer, select a relationship definition and then select **Tools > Relationship Manager** from the menu bar

Relationship Manager starts. At this point it is disconnected from the server; you must connect to an InterChange Server instance as described in "Connecting to InterChange Server" on page 84 to proceed further.

Figure 38 shows Relationship Manager in a disconnected state.




Figure 38. Relationship Manager

## Connecting to and disconnecting from a server

You must connect Relationship Manager to InterChange Server to work with relationship instances and data. Follow the instructions in the following sections to connect Relationship Manager to a server and disconnect Relationship Manager from it:

- “Connecting to InterChange Server”
- “Disconnecting from InterChange Server” on page 85

### Connecting to InterChange Server

Do the following to connect Relationship Manager to InterChange Server:

1. Select **Server > Connect** from the menu bar of Relationship Manager. The Connect to InterChange Server dialog appears (see Figure 39).



Figure 39. Connect to InterChange Server

2. Do one of the following to populate the name of the InterChange Server instance to which you want to connect in the **Server Name** field:
  - Type the name of the InterChange Server instance in the **Server Name** field.

**Important:** The name of an InterChange Server instance is case-sensitive, so you must be sure to be accurate when specifying the name.

  - Select a cached server name from the drop-down menu.
  - Do the following to browse for the InterChange Server instance on the network:
    - a. Click the browse button.
    - b. At the “Get an active server” dialog, select the desired InterChange Server instance from the list.
    - c. Click **OK**.
3. Type the user name to interact with the InterChange Server instance in the **User Name** field.
4. Type the password for the user name supplied in step 3 in the **Password** field.
5. If you do not want to have to supply the user name and password each time you have to connect to the InterChange Server instance in System Manager then enable the **Save user name and password** checkbox.
6. If you want to open a relationship at this time type the name of the relationship definition in the **Relationship** field.

If you do not want to open a relationship at this time you can open it after connecting to the server. For more information, see “Opening a relationship.”

7. Click **Connect**.

If you connect to InterChange Server in Relationship Manager and specify a relationship to open as described in step 6 on page 84, then Relationship Manager displays the Retrieve Relationship Instances window, described in “Retrieving relationship instances” on page 86.

### **Disconnecting from InterChange Server**

To disconnect Relationship Manager from InterChange Server, select **Server > Disconnect** from the menu bar of Relationship Manager.

## **Working with relationships**

Once you have started Relationship Manager and connected it to an InterChange Server, you can use Relationship Manager to work with relationship data as described in the following sections:

“Opening a relationship”

“Retrieving relationship instances” on page 86

“Creating relationship instances” on page 88

“Deleting relationship instances” on page 90

“Deactivating and activating participants” on page 90

“Copying participants” on page 91

“Loading and unloading business object files” on page 91

### **Opening a relationship**

Do the following to open a relationship definition in Relationship Manager after it is already connected to the server:

1. Select **File > Open** from the menu bar of Relationship Manager.
2. At the “Open Relationship” window, select the name of the relationship you want to open.

Figure 40 on page 86 shows the “Open Relationship” window.

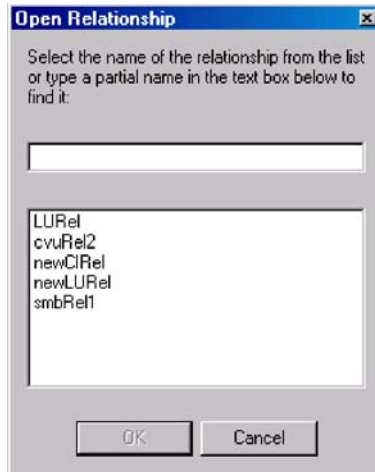


Figure 40. Opening a relationship

3. Click **OK**.

When you open a relationship, Relationship Manager displays the Retrieve Relationship Instances window, described in “Retrieving relationship instances.”

### Retrieving relationship instances

Select **File > Retrieve** from the menu bar of Relationship Manager to retrieve relationship instances or return a count of how many instances there are for a relationship. Figure 41 on page 87 shows the Retrieve Relationship Instances window.



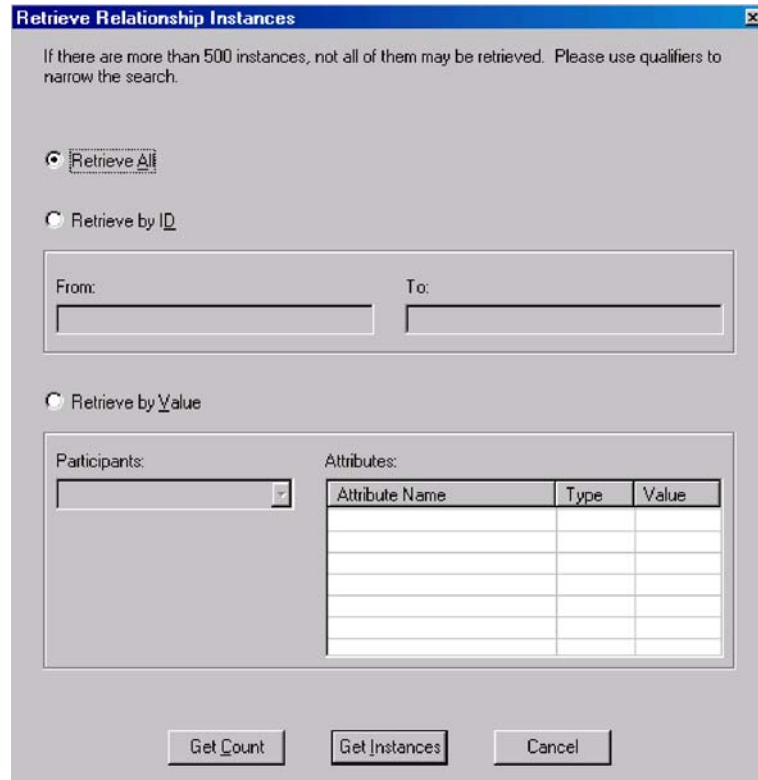


Figure 41. Retrieving relationship instances

The Retrieve Relationship Instances window is also displayed when you specify a relationship to open when connecting to InterChange Server.

You can perform the following operations with the Retrieve Relationship Instances window:

- Retrieve the first 500 instances for the relationship, as described in “Retrieving all instances.”
- Retrieve a range of instances for the relationship based on the relationship instance ids, as described in “Retrieving by relationship ID.”
- Retrieve a relationship instance that contains a participant of a particular value you specify, as described in “Retrieving by participant data” on page 88.
- Return a count of the number of instances for the relationship, as described in “Returning a count of the relationship instances” on page 88.

Depending on the number of participants in the relationship definition and the number of participant instances in each relationship instance, these retrieval queries may take some time.

**Retrieving all instances:** Do the following to retrieve the first 500 instances for a relationship:

1. At the Retrieve Relationship Instances window click **Retrieve All**.
2. Click **Get Instances**.  
Relationship Manager displays the first 500 instances for the relationship.

**Retrieving by relationship ID:** Do the following to retrieve a range of up to 500 instances:

1. At the Retrieve Relationship Instances window click **Retrieve by ID**.

2. Type the ID of the first instance in the range you want to retrieve in the **From** field.
3. Type the ID of the last instance in the range you want to retrieve in the **To** field.
4. Click **Get Instances**.

Relationship Manager displays up to 500 instances in a range of the ids you specify.

**Retrieving by participant data:** Do the following to retrieve a relationship instance based on values for key or non-key attributes of selected participants:

1. At the Retrieve Relationship Instances window click **Retrieve by Value**.
2. Select the participant whose value you want to search on from the **Participants** drop-down menu.

For identity relationships, the drop-down menu lists the participant names followed by the business object definition with which the participant is associated.

For lookup relationships, the drop-down menu lists the participant names followed by the word "Data".

3. Type one of the types of values listed in Table 16 in the **Value** column in the "Attributes" pane.

*Table 16. Supported values for retrieving relationship instances by participant data*

Value	Description
Participant data	The data of the selected participant.  For example, if the relationship is an identity relationship you would specify the ID of the participant instance that you know to find the relationship instance in which it exists.  If the relationship is a lookup relationship you would specify the non-key data value of the participant instance.
%	Any string of characters. This option is case-sensitive; numbers are included in the character set.  For example, if %A were specified for a participant that stores abbreviated forms of the names of the United States, the values CA, GA, IA, LA, MA, PA, VA, and WA would be returned.
_	Any single character.  As an example, _00 would retrieve 100, 200, a00, b00, and so forth.

4. Click **Get Instances**.  
Relationship Manager displays the first 500 relationship instances that match the specified value.

**Returning a count of the relationship instances:** To return the number of relationship instances that satisfy a retrieval criteria, select the options for the criteria as described in "Retrieving all instances" on page 87, "Retrieving by relationship ID" on page 87, or "Retrieving by participant data," and then click **Get Count** instead of **Get Instances**.

### Creating relationship instances

Do the following to create a new instance for a relationship:

1. Create the new relationship instance by performing any of the following tasks:

- Select **File > New Instances** from the menu bar.
- Use the keyboard shortcut **Ctrl+N**.
- Click the **New Relationship Instance** icon in the toolbar.

Relationship Manager displays the new relationship instance.

Highlighted at the top of the hierarchal relationship tree, on the entry line with the relationship icon is the placeholder for the relationship instance ID, which displays three question marks (??). Once you save the relationship instance or any of its participants, InterChange Server automatically generates the new relationship instance ID and Relationship Manager replaces the question marks with this instance ID.

2. Expand the new relationship instance by clicking on the plus (+) sign next to the ??? placeholder icon.

The relationship tree displays participant definitions, participant instances, and participant key and non-key attributes beneath the relationship instance in descending order.

3. Do the following to create a new participant instance in the relationship instance:
  - a. In the relationship tree, select the participant definition for which you want to create an instance.
  - b. Do one of the following to add an instance for the participant:
    - Right-click a participant definition in the listing and choose **Add Participant** from the context menu.
    - Click **Add Participant** in the standard toolbar.
  - c. Expand the new participant instance by clicking on the plus (+) sign next to it.
  - d. Select the new participant instance.
  - e. Click the **Value** column for the participant instance once, then type the desired value into the cell.

**Note:** If the **Value** field for the attribute displays three question marks (??), the participant is managed by InterChange Server. You cannot enter values for these participants because InterChange Server automatically generates them when you save the relationship instance. The value is the same value as the relationship instance ID.

At this point, you can perform any of the tasks in the following table.

*Table 17. Tasks for Participant Data*

<b>Task</b>	<b>Action</b>
Save the participant instance.	To save the new participant instance, right-click the participant instance and choose <b>Save Participant</b> from the context menu. Relationship Manager saves in the appropriate relationship table the data for this participant. The <b>Modified</b> column for the participant instances displays the date the participant was saved, which is the create date, in this case. <b>Note:</b> Once the participant data has been saved, it cannot be changed. To change its data, the participant must be deleted and another created.
Add more participant instances.	Repeat repeat step 3 in the previous list. <b>Note:</b> If you are working with an identity relationship, you cannot create more than one participant instance for a participant definition.

Table 17. Tasks for Participant Data (continued)

Task	Action
Delete a participant.	If necessary, you can delete a saved participant instance by right-clicking on the participant instance and choosing <b>Delete Participant</b> from the context menu. Relationship Manager removes the participant instance from the relationship table. If you do not want to remove the participant instance from the database, use the <b>Deactivate Participant</b> option (see “Deactivating and activating participants” on page 90). A deactivated participant retains its instance ID and its values.
Save the relationship instance.	Save the relationship instance by performing one of the following tasks: <ul style="list-style-type: none"> <li>• Select <b>File &gt; Save</b> from the menu bar (activated when a relationship instance is selected).</li> <li>• Right-click the relationship instance and choose <b>Save Relationship</b> from the context menu.</li> </ul> <p>InterChange Server generates the relationship instance ID and Relationship Manager replaces the ??? placeholder with this new ID. Relationship Manager updates the modified date on all saved participant instances to this date.  <b>Note:</b> At least one participant instance and all key attribute data must be created before the relationship instance can be saved.</p>
Save all relationship instances.	Select <b>File &gt; Save All</b> from the menu bar. InterChange Server generates the relationship instance IDs for any relationship instances that do not have one. Relationship Manager replaces any “???” placeholders with the new IDs. Relationship Manager updates the modified date on all saved participant instances to this date.

### Deleting relationship instances

To delete a relationship instance from the relationship tables, select the relationship instance you want to delete and perform one of the following actions:

- Select **File > Delete** from the menu bar.
- Right-click the relationship instance and select **Delete Relationship Instance** from the context menu.

The relationship instance and its data are deleted from the relationship tables for the current relationship.

### Deactivating and activating participants

A participant instance can be deactivated, or made inactive. Deactivating a participant instance removes it from the relationship instance and prevents it from displaying in the Relationship Manager window, but its record remains in the relationship table so it can be re-activated in the future.

**Deactivating a participant:** To deactivate a participant instance, right-click the participant instance you want to deactivate and choose **Deactivate Participant** from the context menu. The participant is removed from the Relationship Manager display but not from the relationship tables.

**Activating a participant:** To activate a participant instance, take the following steps:

1. Select **View > Show Deactivated Participants** from the menu bar.

The Deactivated Participants window displays as shown in Figure 42 on page 91.

Relation/Participant	Type	Value	Modified	Action
[3]				
partDef2	CustomerB			
[0]			October 2, 2003 11:09:18...	Deactivate
Order_1[0].OrderName	String	dog		
AccNumber	Integer	103		
[4]				
partDef2	CustomerB			
[0]			October 2, 2003 11:09:18...	Deactivate
Order_1[0].OrderName	String	lion		
AccNumber	Integer	105		
[14]				
partDef2	CustomerB			
[0]			October 2, 2003 11:09:13...	Deactivate
Order_1[0].OrderName	String	fish		
AccNumber	Integer	114		

Figure 42. Deactivating participants

2. Select the relationship instance that contains the deactivated participant you want to activate from the list.
3. Expand this relationship instance until the deactivated participant instances display in the list.
4. Right-click the participant instance to reactive and choose **Activate** from the context menu.
5. Select **Edit > Refresh** from the menu bar.

The activated participant instance displays in its relationship instance in the Relationship Manager window.

**Note:** If a participant instance in an identity relationship is deactivated and another participant is added in its place (that is, assigned the same instance ID), the original participant is removed from the Deactivated Participants listing, but remains in the database.

## Copying participants

You can create a new participant instance by copying an existing participant instance. To copy a participant instance, take the following steps:

1. In the relationship instance, right-click the participant definition and choose **Add Participant** from the context menu.
2. Right-click the participant instance you want to copy and choose **Copy Participant** from the context menu.
3. Right-click the newly created participant instance and choose **Paste Participant** from the context menu.

## Loading and unloading business object files

You can load a business object file of the same type into a participant. To load a business object data file into a participant, take the following steps:

1. Right-click the participant instance where you want to load the business object file and choose **Load Participant with Business Object**.

The Participant window displays the business object associated with that participant instance, as shown in Figure 43 on page 92.

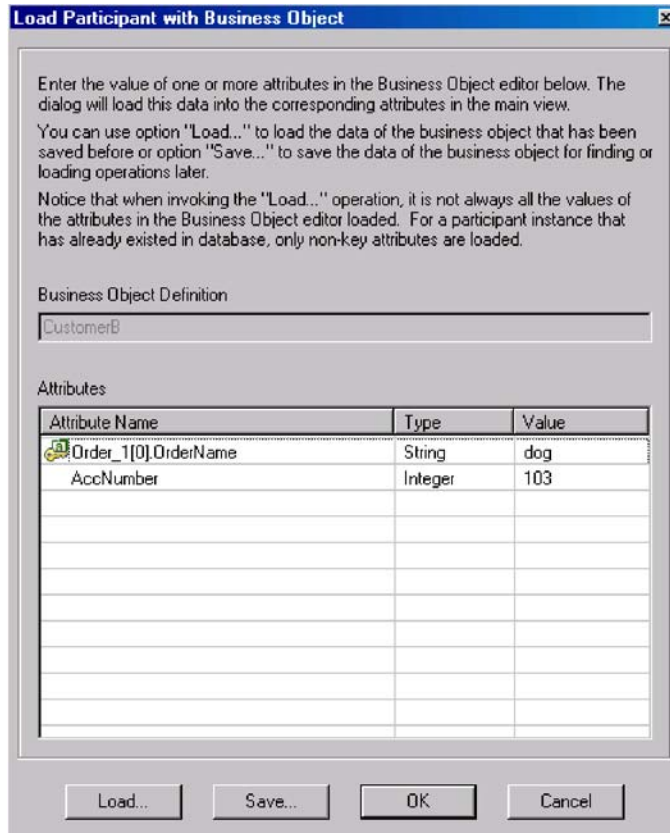


Figure 43. Loading participants with business objects

2. Click **Load**.
3. Navigate to and open the business object file you want to load.
4. Click **OK**.

**Note:** Only the first instance of a relationship is loaded if more than one instance exists in the file.

## Working with relationship data

An important feature of Relationship Manager is its ability to access and manipulate relationship runtime data contained in the relationship tables. The following topics describe how to use Relationship Manager to manipulate and access runtime data:

“Searching for participants”

“Filtering the displayed participants” on page 94

“Cleaning up participants” on page 95

“Printing relationship data” on page 96

### Searching for participants

You can search for participant instances based on different criteria. Depending on how specific your search criteria is, your searches can locate a unique participant instance or a group of participant instances. You can find participant instances in the following ways:

**Finding instances by business object:** This option searches for instances whose data type is an attribute in a business object.

Do the following to search for instances by business object:

1. Select a participant instance in Relationship Manager.
2. Select **Edit > Find Instances by Business Object** from the menu bar.

The Find Instances by Business Object window appears (see Figure 44).

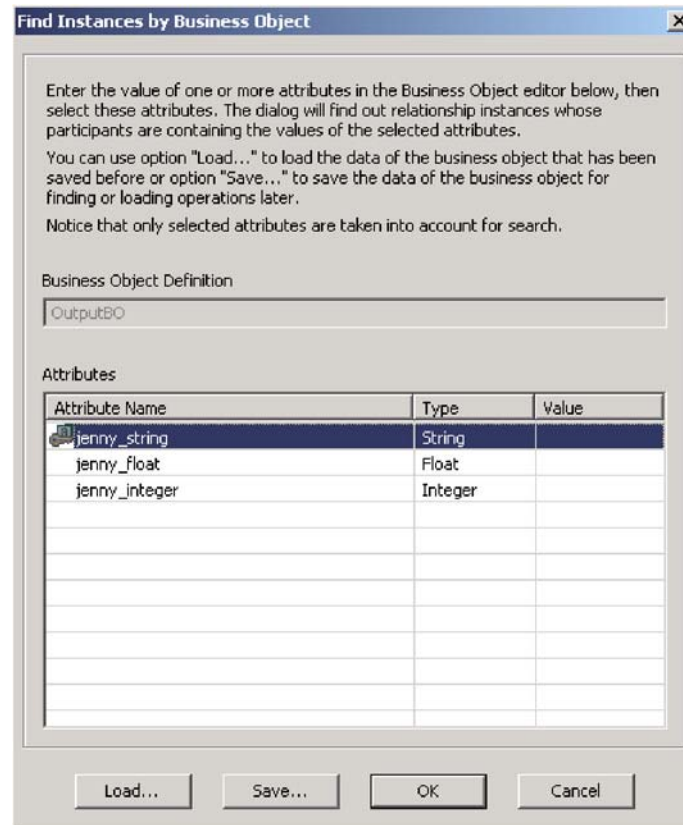


Figure 44. Find Instances by Business Object

3. Type the participant value by which you want to search in the **Value** cell.
4. Click **OK**.

Relationship Manager displays any matching instances in a dialog box.

5. Double-click any of the instances in the dialog displayed by Relationship Manager to navigate to and highlight the instance.

**Finding instances by data:** This option searches for instances whose type is Data.

Do the following to search for instances by data:

1. Select a participant instance in Relationship Manager.
2. Select **Edit > Find Instances by Data** from the menu bar.

Relationship Manager displays the “Find Instances by Data” window, as shown in Figure 45 on page 94.



Figure 45. Finding instances by data

3. Type the participant value by which you want to search in the **Enter the data** cell.
4. Click **OK**.  
Relationship Manager displays any matching instances in a dialog box.
5. Double-click any of the instances in the dialog displayed by Relationship Manager to navigate to and highlight the instance.

### Filtering the displayed participants

You can filter the participants to only display those created or modified between certain dates. Do the following to filter the displayed participants:

1. Select a participant in Relationship Manager.
2. Select **View > Filter** from the menu bar.

Relationship Manager displays the “Filter” dialog, as shown in Figure 46.

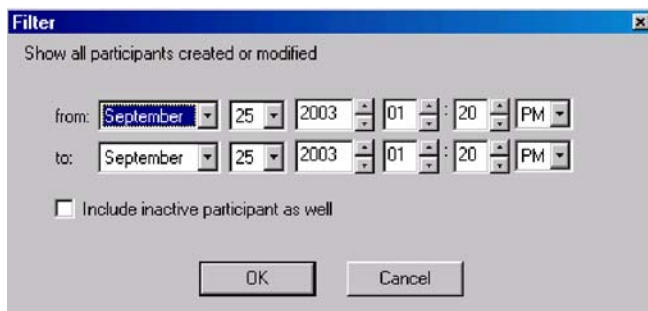


Figure 46. Filtering participant results

3. In the “Filter” dialog, enter the earliest date of creation or modification for the participant in the **from** field and the latest date of creation or modification for the participant in the **to** field.  
Use the following techniques to enter the date value:
  - Type letters in the text field to cycle through days of the week. For instance, type **S** to cycle through dates that occur on Saturdays and Sundays.
  - Click the small up and down arrows to increment or decrement the date by one day.
  - Click the large down arrow to display a calendar that you can use to select a date.
4. Click the **Include inactive participants as well** checkbox if you want to include inactive participants in the resulting display.
5. Click **OK**.



Relationship Manager displays the history of activity for the filtered interval in the “Filter Results” dialog. The filtered display includes inactive participants if the **Include inactive participants as well** option box is checked. Figure 47 shows the “Filter Results” dialog.

Relation/Participant	Type	Value	Modified	Action
[3] partDef2	Custom...			
[0] Order_1[0].OrderName	String	dog	August 8, 2003 7:1...	Create
AccNumber	Integer	104		
[1] Order_1[0].OrderName	String	dog	August 14, 2003 1:...	Update
AccNumber	Integer	104		
[2] Order_1[0].OrderName	String	dog	August 18, 2003 10:...	Activate
AccNumber	Integer	103		
[3] Order_1[0].OrderName	String	dog	August 20, 2003 4:...	Update
AccNumber	Integer	103		
[4]				
[5]				

Figure 47. Viewing filtered participant data

### Cleaning up participants

Do the following to clean up participants due to inconsistent or corrupt data in the source application or generic object:

1. Select a participant in Relationship Manager.
2. Select **Edit > Clean Up Participants** from the menu bar.

Relationship Manager displays the “Clean up Participants” dialog, as shown in Figure 48.

Clean Up Participants

This clean-up will affect all changes to all participants in the current relationship. These changes cannot be undone.  
Below, enter a time to which you want to revert the values of the participants and click OK.

Clean up from: September 25, 2003 01:30 PM

OK Cancel

Figure 48. Cleaning up participants

3. In the “Clean Up Participants” dialog box, enter the date to which you want to revert the participant values to in the **Clean up from** field.

Use the following techniques to enter the date value:

- Type letters in the text field to cycle through days of the week. For instance, type **S** to cycle through dates that occur on Saturdays and Sundays.
- Click the small up and down arrows to increment or decrement the date by one day.
- Click the large down arrow to display a calendar that you can use to select a date.

4. Click **OK**.

All participant adds, deactivations, and activations since that point in time are erased from the database. A participant that has been deleted or whose value has been modified cannot be cleaned up.

### **Printing relationship data**

Relationship Manager allows you to print information about a relationship's runtime data. It creates a tree representation of the runtime data, much like the data appears in the tool's main window.

**Sending relationship data to a printer:** The printing ability of Relationship Manager sends the current contents of the relationship tree in the main window to the printer. Do the following to print relationship runtime data:

1. Expand the relationship tree of Relationship Manager so that the information you want to print is displayed.
2. If you want to print only a portion of the relationship instances, select only those instances by highlighting them.
3. Print relationship runtime data in any of the following ways:
  - Select **File > Print** from the menu bar.
  - Use the keyboard shortcut **Ctrl+P**.
  - Click **Print** in the toolbar.
4. The Print Relationship Instances dialog appears. Select either "all instances" or "selected instances," then click OK.

---

## **Backing up IBM WebSphere InterChange Server components**

Backing up the IBM WebSphere ICS system is among the more critical tasks for IBM WebSphere system administrators. Standardized backup procedures allow for easier environment restoration in the event of system failures. Backing up the IBM WebSphere ICS system is also important because hardware or software failures may leave data in an inconsistent state between IBM WebSphere ICS and the integrated applications.

This section covers the following topics:

"Planning a backup schedule" on page 96

"Backing up components" on page 97

### **Planning a backup schedule**

Plan and carry out procedures for regularly scheduled backups of the IBM WebSphere ICS system. The more frequently you perform backups, the less data you need to recover in the event of data loss. When planning a backup schedule, keep in mind the information in this section.

Within the IBM WebSphere ICS system, two types of data should be backed up: static data and dynamic data.

- Static data rarely changes and should be backed up only when changed. For example, static configuration data stored in the IBM WebSphere ICS repository needs backing up only when it has changed. Static data should also be backed up before any planned reinstallations or upgrades to the system. Following is a partial list of static data in the IBM WebSphere ICS system:
  - IBM WebSphere ICS repository (except for relationship tables)

- Custom collaborations components, such as Java class files (.class), and message files (.msg)
- Custom connectors
- Map components, including: map design files and Java class files (.class).
- Dynamic data constantly changes and should be backed up on a regular basis. For example, the relationship tables maintain the instance data for relationship definitions. Since relationship instance data is maintained continuously, regularly back up this data as well as application data.

The relationship tables are stored by default in the repository database. If you store them in another database, you need to back up that database. For more information about settings for storage of relationship tables, see the *Map Development Guide*.

Following is a partial list of dynamic data in the IBM WebSphere ICS system:

- IBM WebSphere ICS cross-reference database
- IBM WebSphere ICS relationship tables
- IBM WebSphere ICS WIP (event management) and transaction tables
- WebSphere MQ queue data
- IBM WebSphere ICS connector archive tables (this is part of the application backup; all events since the last backup should be archived)
- Log files (as desired for historical information)

Plan your backup schedule at times when your systems environment is in a quiescent state or in a state with a minimal amount of event processing. IBM WebSphere ICS is in a quiescent state when all of the following conditions exist:

- All working queues are drained.
- All collaborations are paused so that no new data can be written to the cross-reference tables.
- All data is consistent between the integrated applications.

## Backing up components

Different components of the IBM WebSphere ICS environment require different backup procedures. The following topics are described in this section:

“Backing up relationship tables” on page 98

“Backing up and loading the repository” on page 98

“Backing up system installation files” on page 98

“Backing up collaboration class files” on page 99

“Backing up archive tables” on page 99

**Attention:** When backing up IBM WebSphere InterChange Server Components, do not back up the WebSphere MQ queues. WebSphere MQ queues represent in-progress transactions in the system, which are dynamic and therefore should never be backed up. Instead, IBM WebSphere ICS recommends that the WebSphere MQ queues be mirrored in a fail-over scenario.

## Backing up relationship tables

Relationship tables are backed up using the standard backup utility for the database where these tables reside. Schedule this backup to coincide with the corresponding application backups. If you back up applications at different times, back up the relationship tables each time you back up an application. There are often static relationship tables within the relationship database. Although this data is static, it is recommended that you back up *all* relationship tables together. Make sure the IBM WebSphere ICS system is in a quiescent state when backing up the relationship tables. For more information on bringing the system to a quiescent state, see “Shutting down InterChange Server” on page 49.

It is recommended that the relationship database log be mirrored to assist in recovery. If hardware/software cost is not a consideration, the relationship runtime data can also be mirrored.

The set of relationship tables for one relationship are closely associated, so you should back up all of these at the same time.

Back up relationship information using the standard backup utility from the DBMS (Database Management System) where these tables reside.

**Note:** To avoid data loss, run relationship backups at the same time you run backups for the applications that the tables reflect.

## Backing up and loading the repository

Repository tables are backed up using the `repos_copy` command. For more information on this command, see “Using Repos\_Copy” on page 99. Back up the repository whenever it is modified and before and after performing a reinstallation or an IBM WebSphere ICS software upgrade. The IBM WebSphere ICS system does not need to be in a quiescent state when backing up the repository.

The method to use for backing up the repository depends on whether your database is partitioned or unpartitioned.

**Backing up partitioned databases:** If your databases are partitioned, you can use the standard database backup utility from the DBMS to back up the Repository, Event Management, and Transaction databases.

**Note:** It is recommended that the Repository, Event Management, and Transaction database logs be mirrored to assist in recovery.

**Backing up an unpartitioned (single) database:** If your IBM WebSphere ICS databases are not partitioned, meaning they are contained in a single database, they should not be part of your normal database backup routine. The IBM WebSphere ICS databases contain transient data whose recovery can cause inconsistencies in the system. Instead, back up the objects in the IBM WebSphere ICS repository by using the `repos_copy` utility.

## Backing up system installation files

The system installation files should be backed up at the following stages:

- After initial installation.
- Periodically during the development phase:
  - After collaboration design and development
  - After connector design and development
  - After map development and customization

- After the configuration and customization phase is complete.

### Backing up collaboration class files

Back up collaboration class files with your other non-IBM WebSphere ICS system files. Coordinate the repository backup with the collaboration class file backups.

### Backing up archive tables

Some applications have archive tables. Back up archive tables using the standard database utility for the database in which they reside. The archive tables are part of the IBM WebSphere ICS system, but typically reside in the application's database. Back up the archive tables on a regular basis. Data in the archive table represents all of the events that have passed from the application to the IBM WebSphere ICS system. These events can be used to "resynchronize" the application and the IBM WebSphere ICS cross-reference tables. For more information, see "Maintaining event archives" on page 145.

---

## Using Repos\_Copy

Repos\_copy is a command line interface for working with integration components and InterChange Server repositories. It allows you to deploy a package—a collection of integration components—to a server repository, or to export components from the repository to a package.

To run repos\_copy, enter the command at a shell prompt (UNIX) or in an MS-DOS command prompt window (Windows). The *ProductDir/bin* directory, where the utility resides, should be in your path as a result of installation.

**Note:** The repos\_copy output file contains encrypted passwords for relationships and connector applications. If you try to edit the output file and change these passwords, repos\_copy will not work.

**Important:** When repos\_copy deploys components to the repository, it deploys them to the repository *only*. It does *not* deploy them to any in-memory tables of business object definitions. For instance, connectors load business object definitions from the repository into their memory space when they start. If you deploy a business object definition to the repository to update it, you must restart the connector agent so that it loads the modified business object definition into memory. You must therefore stop and restart InterChange Server and components that load definitions into memory for them to load recently deployed components.

This chapter has the following sections:

- "Repos\_copy syntax" on page 99
- "Repos\_copy usage scenarios" on page 104
- "Locale for repos\_copy files" on page 110

For more information about backing up the system, see the *System Administration Guide*.

## Repos\_copy syntax

Table 18 on page 100 describes the options of repos\_copy and their arguments, and shows the correct case usage for the options and the lack of spacing between the option and its argument. The syntax shows that the options between curly braces ({} ) represent a set of options that are required. If you do not specify the -u, -p, -i,

-o, or -s options at the command line, then repos\_copy prompts you for them. If you do not specify them when prompted, repos\_copy does not execute. Options enclosed in brackets ([]) are optional.

**Note:** Some new arguments have been added in release 4.2, and some arguments from the previous release have been removed. For a list of these arguments, see “New arguments in release 4.2” on page 104 and “Arguments removed in release 4.2” on page 104.

```
repos_copy [-sserverName] [-username] [-ppassword]
{-i [filename1] [-rrelationshipName [relationshipName2]] [[-k] [-ai] [-ar] [-arp]
[-xcompilePackage] [-vp] [-vr]}
{-o [outfile] [-fEntityFile] [-eEntityType:Entity1 [+EntityType:Entity2] [+...]]
[-deep] [-summary]}
{[-d] | [-doEntityType:Entity [+EntityType:Entity2] [+...]] |
[-dfoEntityType:Entity [+EntityType:Entity2] [+...]]}
{-v}
{-vr}
{[-xCompileAll] | [-xCompileAllCollabs] | [-xCompileAllMaps] |
[-xCompileCollab:collabTemplateName [+collabTemplateName] [+...]] |
[-xCompileMap:nativeMapName [+nativeMapName] [+...]]}
```

Table 18. Repos\_copy command options

Option	Description
-ai	Ignore and do not load any duplicate objects (business objects, maps, relationships, collaboration templates and objects, and connectors) that are found when deploying a package.
-ar	Replace any duplicate objects (business objects, maps, relationships, collaboration templates and objects, and connectors) that are found when deploying a package. <b>Note:</b> The -ar option only works with release 4.2.0 or later.
-arp	This is an interactive version of the -ar option. If the components in the package being deployed already exist in the repository then repos_copy displays a prompt asking if you want to ignore or replace the component. <b>Note:</b> The -arp option only works with release 4.2.0 or later.
-d	Deletes the components in the repository, except the state data. Use this option to delete all of the components from the repository.
-deep	Used with the -e option when you want to include all the dependent components. If you omit the -deep option, only the component that is specified with the -e option will be included.
-dfoEntityType:Entity [+EntityType:Entity2]	This option is the same as the -do option except that it will forcefully delete the component even if the component has referents that depend on it. This option only works with the repository of a server that is running in design mode. A server that is running in production mode does not permit unresolved dependencies and references.

Table 18. *Repos\_copy* command options (continued)

Option	Description
<code>-doEntityType:Entity[+EntityType:Entity2]</code>	<p>Specifies the entities to be deleted from the repository. See Table 19 on page 104 for the list of entity types and keywords. If the object has no referents—other components that depend on it—then the deletion takes place. If the object has referents, then the deletion fails and a message is displayed. The behavior is the same in both design mode and production mode. For more information about starting the server in design mode or production mode, see the <i>System Implementation Guide</i>.</p>
<code>-eEntityType:Entity1[+EntityType:Entity2...]</code>	<p>Exports one or more referenced first-class entities. A first-class entity is a business object, collaboration object, collaboration template, connector, database connection pool, map, or relationship. You identify the entity to load or unload by specifying one of the keywords in Table 19 on page 104.</p> <p>Follow the <i>EntityType</i> keyword with a colon (:) and the name of the entity. Use the “+” to specify more than one entity. When combined with the <code>-o</code> option, the <code>-e</code> option unloads the data to an output file.</p>
<code>-fentityFile</code>	<p>This option is similar to the <code>-e</code> option except that the names of the entities to be imported are stored in a file. The file should contain references to the entities, with the following conditions:</p> <ul style="list-style-type: none"> <li>• The entity names must follow after the proper entity type keyword. The entity types and their keywords are listed in Table 19 on page 104.</li> <li>• A colon must separate the entity type from the entity name.</li> <li>• There must be a new line separating each entity reference.</li> </ul> <p>When combined with the <code>-o</code> option, this option exports the components to a package.</p>
<code>-ifilename</code>	<p>Deploys the specified package file to the repository. If you omit the input file name value, the command interactively prompts you to enter the name of the input file. The file can be either a <code>.jar</code> file containing objects in XML format, or a file in text format from a release prior to 4.2.0.</p> <p>The <code>.jar</code> files created by <code>repos_copy</code> or System Manager have a particular structure which must be maintained for any subsequent imports of such a file to be successful. You should not, therefore, ever modify an input file manually.</p>

Table 18. *Repos\_copy* command options (continued)

Option	Description
-k	Overrides the default behavior of <i>repos_copy</i> when it finds a Mercator map in the package file it is loading. By default, <i>repos_copy</i> exits if it encounters a Mercator map. If you use the -k option, <i>repos_copy</i> skips over any Mercator maps in the package file and proceeds with the deployment process.
-mode	Returns the mode of the server. For more information about InterChange Server modes, see the <i>Implementation Guide for WebSphere InterChange Server</i> .
-ncencoding	Specifies the character encoding when importing a text-based repository file from releases prior to 4.2.0.  For a list of valid character encodings, see the Java documentation about the String class.
-outfilename	Exports the components in the repository to the specified package file. You must specify the name of the package file. If the file already exists then <i>repos_copy</i> prompts you to overwrite it or not. The output file is in .jar format, and contains the component definitions in XML format, as well as .java source files for components that have them. This option cannot be combined with the -i or -d options, nor can it export components in text format as it did in previous releases. <i>Repos_copy</i> does not append the .jar extension, so you must specify it when specifying the name of the output file.
-ppassword	Specifies the password for the user name supplied with the -u option. The password case-sensitive. If you do not specify this option then <i>repos_copy</i> prompts you for the password.
-r*	This option is similar to the -r option; it allows you to import relationship definitions and not create the runtime schemas for any of them.
-rrelationshipName1[:relationshipName2]	Loads the named relationship definition(s) into the repository without creating its runtime schema.
-sserverName	Specifies the name of the InterChange Server instance with which <i>repos_copy</i> should interface. The name is case-sensitive. If the server name is not specified, the tool prompts for a server name.
-summary	This option prints a list of components in the server repository (they are identified as “artifacts” rather than as components in the output). The output is in XML format. You can combine this option with the -o option to print the output to a file rather than the console.
-username	Specifies the user name to log in to InterChange Server. If no user name is specified, <i>repos_copy</i> prompts for a user name.



Table 18. *Repos\_copy* command options (continued)

Option	Description
-v	Prints the version number of the program that the <i>repos_copy</i> utility executes.
-vp	This option validates a package file. The server validates packages against the repository and makes sure that the dependencies among the components in the package are resolved. If the validation is not successful, <i>repos_copy</i> prints a list of the missing dependencies. This option does not make any changes to the repository; it just validates the package file. When using the -vp option you must also use the -i option to specify the package file to be validated.
-vr	This option validates the repository. The output message indicates whether the validation is successful or not. If the validation is not successful, <i>repos_copy</i> prints a list of the missing dependencies.
-wi	When this option is specified, <i>repos_copy</i> does not display any warnings that occur during the compilation of collaboration templates or maps. Only errors that occur during compilation are displayed. This allows the user to ignore warnings about deprecated methods, for instance.
-xCompileAll	Compiles all collaboration templates and maps in the repository. Valid only for collaboration templates and maps created using release 4.2 or later.
-xCompileAllCollabs	Compiles all collaboration templates in the repository. Valid only for templates created using release 4.2 or later.
-xCompileAllMaps	Compiles all maps in the repository. Valid only for maps created using release 4.2 or later.
-xCompileCollab: <i>collabTemplateName</i> [+ <i>collabTemplateName</i> ]	Compiles the specified collaboration templates in the repository. Valid only for templates created using release 4.2 or later.
-xCompileMap: <i>nativeMapName</i> [+ <i>nativeMapName</i> ]	Compiles the specified maps in the repository. Valid only for maps created using release 4.2 or later.
-xCompilePackage	This option automatically compiles the package being deployed to the server. Since the production-mode server automatically compiles all packages, this option applies only to design-mode servers. For a full description of InterChange Server modes, see the System Implementation Guide. <b>Note:</b> This option works only if you are deploying components from release 4.2. If the components are from a prior release, this option will be ignored.

Table 19. Keywords for different entity types

Entity type	Keyword
Business object	BusObj
Collaboration object	Collaboration
Collaboration template	CollabTemplate
Database connection pool	ConnectionPool
Connector	Connector
Map	Map
Relationship	Relationship

## New arguments in release 4.2

The following list contains all of the new options provided with `repos_copy` in release 4.2:

- `-dfoEntityType:Entity[+EntityType:Entity2]`
- `-doEntityType:Entity[+EntityType:Entity2]`
- `-mode`
- `-r*`
- `-summary`
- `-wi`
- `-xCompilePackage`

## Arguments removed in release 4.2

The following table lists the repos copy arguments that were removed from the

Table 20. Arguments removed from repos copy

Arguments removed	Reason for removal
<code>[-xCompileUpdated]</code> <code>[-xCompileUpdatedCollabs]</code> <code>[-xCompileUpdatedMaps]</code>	All compile update options have been removed, because the server does not support maps or templates prior to release 4.x.
<code>[-xUncompress]</code>	In the new package format, all definitions are stored in the Java Archive (JAR) format, instead of with a proprietary compression algorithm. The <code>-xUncompress</code> argument is therefore no longer necessary.
<code>[-eProject][-w]</code>	Project type is no longer supported for the <code>-e</code> option. Projects are now maintained in the local file by System Manager, instead of being maintained by the server.

## Repos\_copy usage scenarios

This section describes many of the common situations in which you will use `repos_copy`. It contains the following sections:

- “Printing the `repos_copy` command” on page 105
- “Validating a package” on page 105
- “Validating a package” on page 105
- “Deploying a package to the repository” on page 105
- “Validating the repository” on page 107
- “Deleting components from the repository” on page 108
- “Exporting components to a package” on page 109

- “Printing a list of components in the repository” on page 109

## Printing the repos\_copy command

You can run `repos_copy` without any arguments to have the command and its arguments printed out. The example below shows `repos_copy` when executed without any arguments, and the resulting output:

```
C:\>repos_copy
No Command line arguments to ReposCopy were specified
Usage: repos_copy {-o[outputFile] | -i[inputFile]}
      [-sserverName] [-uuserName] [-ppassword]
      [-ai] [-ar] [-arp] [-d] [-k] [-v]
      [-entityType:entityName1[+entityType:entityName2] -deep]
      [-fentityFileName]
      [-rrelationshipName1[:relationshipName2] ]
      [-xCompileAll] [-xCompileAllCollabs] [-xCompileAllMaps]
      [-xCompileCollab:collabTemplateName[+collabTemplateName]]
      [-xCompileMap:nativeMapName[+nativeMapName]]
      [-xcompilepackage]
      [-mode]
      [-doentityType:entityName1[+entityType:entityName2] -deep]
      [-dfoentityType:entityName1[+entityType:entityName2] -deep]
      [-summary]
      [-vp]
      [-vr]
```

## Validating a package

You can validate a package of components before deploying the package to a server. This is very useful because if you deploy a package to a production-mode server all the dependencies must be resolved or the deployment will fail. You cannot validate a user project or integration component library in System Manager to make sure that the dependencies are satisfied, so the only way to find out if a package is valid when deploying with System Manager is to attempt the deployment and use the error information when it fails to resolve the dependencies. If there are many components in the package, this can be a very time-consuming process.

Although you cannot validate an integration component library, you can export it to a package file and then validate the package file using `repos_copy`.

To validate a package file using `repos_copy`, use the `-i` option to specify the name of the package file to be validated and the `-vp` argument to validate it rather than deploy it.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-iWebSphereICS420DEVServer.jar -vp
```

`Repos_copy` validates the contents of the package and displays a message to indicate whether or not the dependencies are resolved.

## Deploying a package to the repository

The `-i` option allows you to deploy a package of components to the repository. If you do not specify the name of the package file then you are prompted to enter it.

The following example shows a file named `WebSphereICS420DEVServer.jar` being deployed to a repository:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-iWebSphereICS420DEVServer.jar
```

**Working with duplicate components during deployment:** Commonly there will be components with the same name in the package file as there are in the

repository. In this case you must decide whether or not you want to replace the components in the repository with those in the package file. The `-ai` option specifies that duplicate components should not be loaded into the repository:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1  
-iCustomer.jar -ai
```

If you want to replace all the duplicate components in the repository, use the `-ar` option as in the following example:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1  
-iCustomerSyncInterface.jar -ar
```

You can use the `-arp` option to interactively replace duplicate components in the repository. This lets you decide for each individual duplicate component whether it should be replaced or not.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1  
-iCustomerSyncInterface.jar -arp
```

**Note:** The `-ar` and `-arp` options only work with release 4.2.0 or later.

**Compiling and creating schemas:** For maps and collaboration to execute at runtime, the maps and collaboration templates defined in the repository must be compiled. For relationships to function properly at runtime, their schemas must be created.

When you deploy components to a server running in production mode, all templates are automatically compiled and all relationship schemas are created. For the deployment to succeed, then, the code of the map and collaboration templates must be valid and InterChange Server must be able to communicate with the databases specified in the settings of the relationship definitions.

When you deploy components to a server running in design mode, the templates are not automatically compiled; relationship schemas are automatically created. There are options you can use to compile the templates, however, and there are options to not create relationship schemas.

The following example uses the `-xCompilePackage` option and does not use any form of the `-r` option. The result is that when the package specified by the `-i` option is deployed, the maps and collaboration templates are compiled and schemas are created for the relationships:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1  
-iWebSphereICS420DEVServer.jar -xCompilePackage
```

You may not want relationship schemas created when you do a deployment. For instance, if you are deploying a package from one environment to another and did not change the properties of the relationships to use the database resources in the new environment then you will not want the schemas created until after you have changed the relevant properties. The following example uses the `-r*` option to not create schemas for all of the relationships in the package being deployed:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1  
-iWebSphereICS420DEVServer.jar -xCompilePackage -r*
```

**Note:** You can use the `-r` option without the asterisk to specify the names of individual relationships whose schemas should not be created. For instance, `-rCustomer:Order` would not create schemas for the Customer and Order relationships, but would still create schemas for any other relationships in the package being deployed.

**Important:** Although there are options to compile maps and collaboration templates after deployment, there is no way to either through `repos_copy` or System Manager to create the schema for a relationship other than during deployment. So, if you chose not to create the schema for a relationship during deployment because you needed to change the database settings, then you need to re-deploy the relationship afterwards and allow `repos_copy` to create the schema for the relationship.

### Validating the repository

The repository must be in a valid state for a server instance to start in production mode. The reason for this is that ultimately the repository must be valid for the server to process flows successfully. Use the `-vr` option to validate a server repository, as in the example below:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull -vr
```

If the server is valid then `repos_copy` writes the following output to the console:  
Validation Succeeded.All Dependencies Resolved.

If the repository is not valid then `repos_copy` prints a list of the dependencies that must be resolved.

### Compiling components in the repository

If you deployed maps or collaboration templates to the repository and did not compile them during deployment, you can use `repos_copy` to compile them afterwards. This can be useful in situations where there are many components to deploy because deployment can take a long time and compiling can make the operation take even longer. Waiting until after the deployment has succeeded to do the compilation task can reduce the risk of spending an even greater amount of time migrating the environment if an error occurs.

The following example shows the use of the `-xCompileAll` option to compile all maps and collaboration templates in the the repository:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-xCompileAll
```

There are options to compile all of either type of component as well. Use `-xCompileAllCollabs` to compile all the collaboration templates, and `-xCompileAllMaps` to compile all the maps. The example below shows the use of `-xCompileAllMaps`:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-xCompileAllMaps
```

Just as you can compile all of one type of component, you can also compile an individual component. Use the `-xCompileCollab` or `-xCompileMap` option followed by a colon and the name of the collaboration template or map to compile a single component. The example below would compile a collaboration template named `CustomerSync`:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-xCompileCollab:CustomerSync
```

**Note:** The compilation options only work with WebSphere InterChange Server version 4.2.0 and later.

## Deleting components from the repository

There are several options provided by `repos_copy` for deleting components in the repository. You can delete the entire repository, individual components, and individual components as well as any components that reference them.

**Note:** Components must be inactive for you to delete them. If you delete a single component then you must deactivate it first or the delete operation will fail. If you want to delete a component and all the components that reference it, you must deactivate not only the single component, but all those that reference it as well. You can delete the entire repository while the components are in an active state. Use System Monitor or web-based System Monitor to manage the states of components. System Monitor and web-based System Monitor are described in the *System Administration Guide*.

**Deleting the entire repository:** Use the `-d` option to delete all of the components in the repository. The following example shows the syntax:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin  
-pnull -d
```

`Repos_copy` presents a prompt asking if you want to delete the entire repository or not.

**Deleting components without referents:** If a component does not have any referents—other components that reference it and require it to exist in order to perform their function in the system—then you can delete the individual component.

Use the `-do` option followed by the entity type, a colon, and the name of the component. The entity types are listed in Table 19 on page 104. The following example deletes the relationship named **Customer**:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin  
-pnull -doRelationship:Customer
```

**Deleting components with referents:** If a component does have referents—other components that reference it and require it to exist in order to perform their function in the system—then you can only delete the component if the server is running in design-mode, and by using certain options.

*Forcing a delete in spite of references:* If a component has referents, `repos_copy` will not let you delete it with the `-do` option. You must use the `-dfo` option to force deletion of a component with referents. Forcing deletion of a component that has referents will leave the repository in an inconsistent state, and a server running in production mode does not permit that, so this option only works with a design-mode server. The following example shows the use of the `-dfo` option to delete the `Order` business object in spite of the fact that other components in the system (such as maps and relationships) have references to it:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-dfoBusObj:Order
```

*Deleting the referents as well:* Another way you can delete a component that has referents is to use the `-deep` option to delete the referents as well. This deletes the component and all of the components that have references to it. The following example shows the use of the `-deep` option when using the `-do` option to delete the `Customer` business object:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-doBusObj:Customer -deep
```

This option, unlike the `-dfo` option, is supported with servers running in production mode because the deletion of the referents along with the component guarantees that the repository remains valid. Keep in mind, however, that it can result in many components being deleted; you should be aware of the implications of this action prior to taking it.

## Exporting components to a package

The `-o` option allows you to export components from the repository to a package. You must specify the name of the package file. When the `-o` option is used alone the entire repository is exported to a file, as in the following example:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-oWebSphereICS420DEVServer.jar
```

You can specify individual components to be exported by using the `-e` option. You must use the `-e` option with the appropriate `EntityType` keyword listed in Table 19 on page 104, and must follow the keyword with the name of the component. You can specify multiple components by concatenating them with the plus (+) sign. In the following example, the Customer business object and CustomerSync collaboration template are exported to a package named `CustomerSyncInterface.jar`.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-eBusObj:Customer+CollabTemplate:CustomerSync -oCustomerSyncInterface.jar
```

You can use the `-deep` option to export the dependencies of a component as well. In the previous example, the Customer business object was exported, but none of its child business objects were. The following example uses the `-deep` option to export the `CustomerSync_ClarifyToSAP` collaboration object and all of its dependencies.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-eCollaboration:CustomerSync_ClarifyToSAP -oCustomerSyncInterface.jar -deep
```

If you want to export specific components, but do not want to have to enter the entity type keyword and component names, you can store them in a text file and use the `-f` option. This is very convenient when you want to frequently export the same components. The following example uses the `-f` option to load the components listed in a text file named `Components.txt` :

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-fComponents.txt -oCustomerSyncInterface.jar -deep
```

The contents of the file `Components.txt` are shown below; a paragraph return follows each entity type keyword and name combination:

```
BusObj:Customer  
Relationship:Customer  
CollabTemplate:CustomerSync
```

**Note:** `Repos_copy` and System Manager are unfortunately inconsistent with respect to what they identify as “dependencies”. If you attempt to delete a component using `repos_copy` but there are components that depend upon it then `repos_copy` lists those referring components as dependencies. However, if you right-click the component in System Manager and select **Show Dependencies** from the context menu the tool lists the components that the selected component depends on.

## Printing a list of components in the repository

You can use the `-summary` argument when executing `repos_copy` to print a list of the components in the repository. The output is presented in XML format.

Although it is not particularly useful to view at the command line, you can combine the `-summary` argument with the `-o` argument to redirect the output to a file and then view the file in a browser or XML editor. The command usage in this case would be the following:

```
C:\>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull -summary -oRepository.xml
```

## Locale for repos\_copy files

The `repos_copy` utility reads metadata from the repository and writes the data out to files in Unicode (UTF-8 format). It also reads such files and loads them into the repository in Unicode (UTF-8 or UCS-2, as the underlying repository database dictates).

`Repos_copy` files created with IBM WebSphere ICS version levels earlier than 4.1.1 can be loaded into the repository correctly only if the dates and times for the component schedules are in full US format. (This is usually not an issue. `Repos_copy` saves all schedule dates in full US format only. The incompatibility could typically arise if the `repos_copy` files have been manually edited.)

---

## Scheduling jobs

Scheduling jobs allows you to create schedules to manipulate the operational states (start, stop, and pause) of connectors and collaborations. By manipulating component states, you can better manage how InterChange Server processes events. You can distribute the server's workload over scheduled time periods, thereby reducing traffic and allowing for more efficient resource management. This section covers the following topics:

“About scheduling jobs” on page 110

“Creating schedules” on page 113

“Modifying schedules” on page 114

“Deleting schedules” on page 115

“Displaying schedules” on page 115

“Enabling or disabling schedules” on page 115

## About scheduling jobs

Scheduling jobs is done through the Schedule window (see Figure 49). From the Schedule window, you can create, modify, and delete scheduled items. You can see a list of all the schedules that are defined for components, or selectively view schedules based on your requirements. You can also enable or disable all schedules on the server.



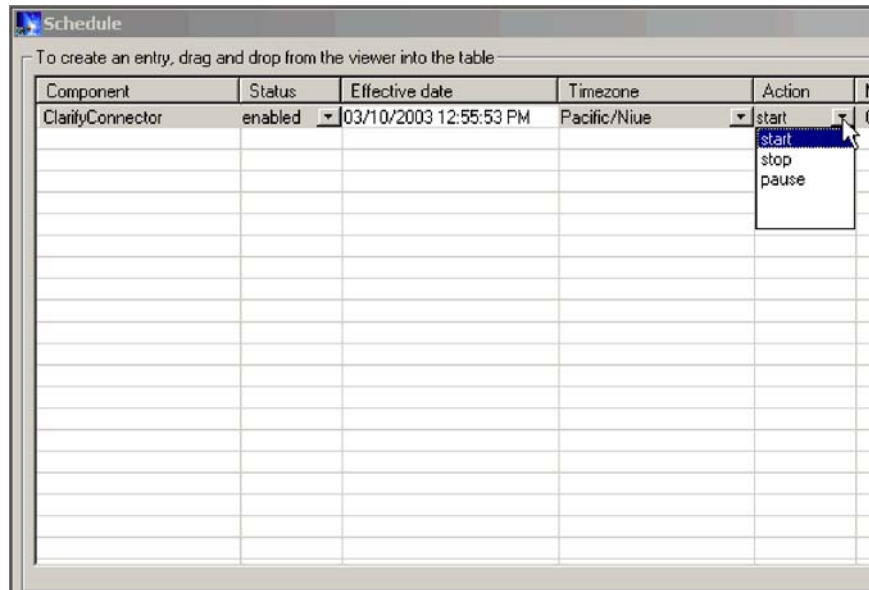


Figure 49. Schedule window

When you create a schedule for a component, you supply information such as when and how often (recurrence) an action (state change) occurs. By default, no schedules are defined for a component. You can define as many schedules as you want for a component. Once a schedule is set, you can enable or disable its use.

The Schedule window allows you to determine the following items:

Status	Enable turns the schedule on and Disable turns the schedule off. The default status is enabled.
Effective Date	The date and time the schedule is enabled. The default is the current date and time.
Timezone	The time zone where the server is located. The default is Pacific Standard time.
Action	The action the schedule performs. Actions are Start, Pause, and Stop.
Next Occurrence	The next time the scheduled action occurs. If the schedule is non-recurring, the date is the same as the Effective Date. If the schedule is disabled, this field is blank.
Component	The name of the connector or collaboration being scheduled.
Comments	Text field that contains comments you enter about the schedule.

If you choose to make the schedule recurring, you can choose from several options including daily, weekly, or monthly.

Because each schedule consists of one action that occurs at a specified time, to create an interval when the server processes a component, you must define both a schedule to start and end processing. As an example, for a connector, you can create one schedule to start processing events at 1 A.M., and another schedule to

pause processing at 3 A.M., daily. Only during that two-hour time period can the connector deliver events to InterChange Server for processing by collaborations that subscribe to that connector.

### **About scheduling connectors**

When you schedule the connector operation, the state you select (start, pause, or stop) determines to what extent work is processed. For example, when you start a connector, it constantly polls an application for new events. When you pause a connector, it stops polling until started again, but is still able to handle service call requests from InterChange Server. A stopped connector is inactive.

By manipulating connector activity with collaboration activity, it is possible to schedule dedicated event processing for an application during a specified time window. To do this, both the collaboration and connector must be running during the same time interval. If the connector was paused, events that were queued can be processed when the connector resumes its activity.

### **About scheduling collaboration objects**

As with connectors, when you schedule the collaboration Object's operation, the state you select (start, pause, or stop) determines to what extent work is processed. To review the collaboration states, see "Collaboration object states" on page 69. For example, when you start a collaboration object, it processes the business objects that it receives from connectors. When you stop a collaboration object, all subsequent events are ignored. So unless you must stop the collaboration object, pause it instead.

**Attention:** Stopping a collaboration object can cause the connector to delete events as unsubscribed. As a warning, the system produces a message if you select Stop.

When you pause a collaboration object, events remain in the collaboration queue until you restart the collaboration object.

**Note:** If a scheduled collaboration object is part of a collaboration group, all collaboration objects in that group are scheduled with the same action.

By manipulating collaboration object activity with connector activity, it is possible to schedule dedicated event processing for an application during a specified time window. To do this, both the collaboration object and connector must be running during the same time interval. By assigning different processing windows to collaboration objects that are bound to the same connector, you can distribute the workload, and to some extent, control the amount of traffic a connector must handle. For example, in Figure 50, each collaboration object gets a dedicated time period when the connector is processing only that collaboration object's events.

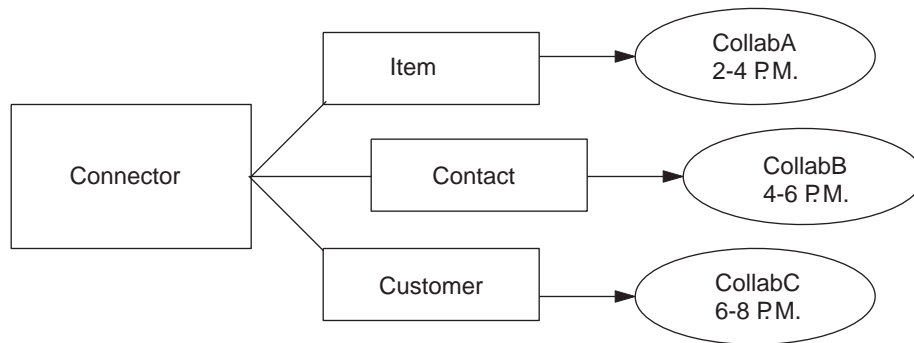


Figure 50. Dedicated processing

Multiple collaboration objects can subscribe to the same business object. In that case, the object is sent to InterChange Server, where it remains until it is picked up by each collaboration object that subscribes to it, when the collaboration object is started after being paused.

### Overriding schedules

Using System Monitor, you can override the state of a scheduled component (for example, start a collaboration object that the scheduler stopped a few minutes ago). Or you can set it to a state to one that the scheduler cannot change. For instance, if a collaboration object is scheduled to pause, you can stop it, not allowing the scheduler to pause it (a collaboration object cannot transition from stop to pause). In such a case, the scheduler does not override the manual change, but logs an error instead.

## Creating schedules

To create a schedule for a collaboration or a connector:

1. Open the Schedule window by right-clicking the Schedule folder in System Manager, then selecting "Edit components' schedule." The Schedule window appears, as shown in Figure 49.

2. From System Manager, select the collaboration object or connector to be scheduled and drag it to the Schedule window.

A new line entry with the name and type of the component is created in the Schedule window (for example, ClarifyConnector (Connector)).

3. Enter information about the schedule by clicking the down arrow in each of the schedule cells:

- a. In the Status field, accept **Enable** to turn the scheduled item on or select **Disable** to turn it off.

An enabled schedule is effective as soon as you click OK or Apply; a disabled schedule is immediately dormant until enabled. When a schedule is disabled, the Next Occurrence cell is blank to indicate there is no scheduled occurrence for this schedule item.

- b. In the Effective Date field, use the calendar to select the date and time when the scheduled item will occur.

By default, the current date and time are set. Use the *MM/DD/YYYY hh:mm:ss* format. A 12 or 24 hour clock is used, based on the Time format configured in the Preferences window, which is available from the Edit menu.

- c. In the Timezone field, select the name of the time zone where the scheduled item is being created, if necessary. By default, the time zone for the scheduled item is set to Pacific Standard Time.

For example, the schedule for a connector is created in New York (select **Eastern Standard** time) while InterChange Server is located in Japan. InterChange Server uses this information to determine the local time for the schedule so it can run the job at the appropriate time.

- d. In the Action field, select the action to be performed. Actions are **Start**, **Pause**, and **Stop**.
  - e. Type in any comments you may have in the Comments text cell. A maximum of 255 characters is allowed.
4. If you want this schedule to be ongoing, click the Recurrence check box and enter information about the next occurrence of the action. Click one of the radio buttons to determine a style for inputting the recurrence information and use the down arrow menus to select specific date information:
    - The first radio button, Every, specifies a number and a date element, such as every 2 days or every 3 weeks.
    - The second radio button specifies the date in terms of a monthly event by the day of the week, such as the first Tuesday of every month or the fourth Friday of every month.
    - The third radio button specifies the date as the last day of some number of months, such as the last day of every 3 months.

If you do not enable the recurrence option, the Next Occurrence field is blank and the schedule expires after it runs. Consistency checks are made to ensure that only one action is scheduled for a particular component on a given date and time. No checks are performed for scheduling conflicts.

**Note:** InterChange Server automatically handles changes between standard and daylight savings time for recurring events.

5. Click either of the Show option check boxes to display specific information about schedules. The Show options are:
  - Show Dependencies, which displays schedules for a Collaboration object's bound connectors and collaborations.
  - Show Expired, which displays schedules that have already processed and whose time to run has expired. Only non-recurring schedules expire.
6. Click OK or Apply to create the scheduled item, which is effective once InterChange Server receives the information.

When InterChange Server and components are geographically distant, there can be a slight delay. If you need to immediately change the state of a component, it is preferable to use System Monitor to start, stop, or pause a component rather than the scheduler.

**Tip:** To schedule a time interval when events are processed for a component, you must create a schedule with the Start action and another with the Stop or Pause action. See "About scheduling jobs" on page 110 for information and examples about determining start and end schedules.

## Modifying schedules

To modify an existing schedule for a collaboration or a connector:

1. Right-click the component in System Manager, then select "Edit Components' schedule." The Schedule window appears (see Figure 49).
2. Edit any field in the Schedule list window to change its value.

To edit Recurrence options, click the cursor anywhere on the scheduled item row; the recurrent values for that scheduled item display in the Recurrence pane if they have been assigned.

3. Click OK to save changes and exit, or click Apply to save changes and keep the window open.

## Deleting schedules

To delete an existing schedule for a collaboration object or a connector:

1. Right-click the component from System Manager, then select "Edit Components' schedule." The Schedule window appears (see Figure 49).
2. Select a scheduled item in the schedule list and click the Delete button (or use the keyboard Delete key) to remove the schedule.
3. Click OK to save changes and exit, or click Apply to save changes and keep the window open.

## Displaying schedules

To display a schedule or a group of schedules:

1. Select and open an object for displaying schedules:
  - Collaboration or connector icon from System Manager. The Schedule window displays all schedules defined for that object. If Show Dependencies is active, all schedules for components connected to the object are displayed.
  - Collaboration or connector folder from System Manager. The Schedule window displays schedules for all objects in that folder.
  - InterChange Server from System Manager. The Schedule window displays all schedules defined for that server, with times shown for the server's time zone.
  - If you choose the Schedule option from the main window, all schedules in the system are displayed.
2. Click any of the column headings to sort schedules by that column.

## Enabling or disabling schedules

To selectively disable or enable schedules:

1. Select an object for displaying schedules.  
See "Displaying schedules" on page 115.
2. Enable or disable the schedule:
  - To enable or disable all schedules, click either the Enable All or Disable All radio button.
  - To enable or disable a single schedule, click the down arrow in the Status column and choose the Enable or Disable option.
3. Click Apply to complete this task.
4. Click OK to exit.



---

## Chapter 3. Troubleshooting the system

This chapter provides troubleshooting topics to help determine and resolve problem scenarios that may occur in an InterChange Server (ICS) system. The following topics are covered:

“Using log and trace files for troubleshooting” on page 117

“Controlling server memory usage” on page 143

“Maintaining event archives” on page 145

“Managing WIP connections” on page 145

“Managing failures” on page 145

“Using Flow Manager” on page 152

“Resolving collaboration deadlocks” on page 162

“Managing WebSphere MQ” on page 165

“Managing database connection pools” on page 166

“Managing High-Availability (HA) systems” on page 167

“Managing runtime properties” on page 169

“Managing the Object Request Broker” on page 170

---

### Using log and trace files for troubleshooting

This section describes how to use log and trace files for troubleshooting. The following topics are covered:

“About logging” on page 117

“Configuring e-mail notification of log messages” on page 123

“About tracing” on page 126

“Configuring logging and tracing” on page 129

#### About logging

Logging is used to communicate system messages, component state changes, failures, and tracing information. Messages that are generated from InterChange Server, collaboration objects, and connectors are sent to the destination you specified when you installed IBM WebSphere ICS, by default, STDOUT (standard output). Messages that are generated from the connector agents are sent to STDOUT, but can be configured to be sent to a separate log file at the agent’s location. The messaging system is always active and provides an accurate monitor of the system.

**Note:** If you configure the messages generated from the connector agent to be sent to a separate log file, you must specify a log file or location that is separate from the InterChange Server log file.

You can configure the messaging system to send messages to a log file or an e-mail recipient in addition to the standard output. You can configure backup files (archives) for the log file, as well as determine their size. Tracing, which is disabled by default because of its drain upon system resources, can be configured when problems arise and detailed information is needed for troubleshooting.

Two tools provide a graphical user interface for configuring and viewing message logging and tracing. Use the:

- The Edit Configuration tool launched from System Manager is used to set up or change system messaging and tracing for InterChange Server, collaborations, and business objects
- Log Viewer is used to display message and trace logs and to search the system message file for an explanation of a message

**Note:** Log Viewer is an IBM WebSphere ICS system tool, which means it runs only on Windows 2000 or Windows XP operating systems. To configure or view a UNIX log file or message using Log Viewer, copy the file from the UNIX machine to a Windows machine that has the IBM WebSphere ICS product installed.

In addition to using Log Viewer to view logs, you can open the log with a text editor or create your own tools to filter the log file.

For information about viewing logging and tracing messages using Log Viewer, see “Viewing log messages” on page 137.

For background information about tracing, see “About tracing” on page 126.

This section covers the following topics:

“Collaboration object messages” on page 118

“Connector messages” on page 119

“Message formats” on page 119

“System logs” on page 120

“Log/Trace file management” on page 121

### **Collaboration object messages**

A collaboration object can generate messages to report runtime information, warnings, and errors. For example, a collaboration might log its decision points and the results of operations.

As InterChange Server executes collaboration objects, it writes their messages to its log. For information about configuring system logging, see “Configuring logging and tracing” on page 129.

In addition, you can send a collaboration object’s messages by e-mail to one or more recipients. You can specify a separate set of e-mail recipients for each



collaboration object. For information on the rules for using e-mail notification, refer to “Configuring e-mail notification of log messages” on page 123.

## Connector messages

Connector messages are sent to the InterChange Server message destination. Depending on your operating system, messages appear in one of the following ways:

### UNIX

A connector logs messages to STDOUT by default, then those messages are rerouted to `connector_manager_<name_of_connector>.log`.

### Windows

A connector logs messages to STDOUT by default, but can be configured to send to a local destination log file or the InterChange Server logging destination.

For information about connector logging, see “Connector Agent logs” on page 121.

To aid in troubleshooting, a temporary log file is created during the connector agent bootup that contains metadata obtained from the connector controller. The metadata consists of business object specifications, properties, and delta-supported properties. The file is named `connectorname.tmp.log` and is found in the `ProductDir\Connectors` directory.

## Message formats

All messages are formatted so they can easily be filtered. Logged messages for InterChange Server and connectors use the same format, which is described in Table 21. When business objects are configured for flow tracing messages, they use these fields and the additional fields, denoted by an asterisk in Table 21. A message delivered to InterChange Server has the following format, using some or all of the following parameters:

*Time: System Name: Thread: MsgType MsgID: SubSystem: FIID: BO: MsgText: BOD:*

Table 21. Message format

Variable	Description
<i>Time</i>	Timestamp: the date of logging in the format <i>year/month/date time</i> .
<i>System</i>	The type of component (system identifier). It can be Server, Collaboration, Business Object, or ConnectorAgent.
<i>Thread</i>	Thread name and thread ID
<i>Name</i>	The name of the component, such as ClarifyConnector.
<i>MsgType</i>	Indicates the severity of the message. See Table 22 on page 120.
<i>MsgID</i>	The message number.
<i>SubSystem*</i>	The subsystem of the current system. It can be Event Management, Messaging, Repository, or Database Connectivity.
<i>FIID*</i>	The flow initiator ID of the business object.
<i>BO*</i>	Business object name.
<i>MsgText</i>	The associated text for the message number.
<i>BOD*</i>	Business object dump. The data contained in the business object.

Following is an example of a message for the server: [Time: 2001/06/07 11:01:29.487] [System: Server] [SS: REPOSITORY] [Thread: VBJ ThreadPool Worker (#-1767149274)] [Type: Trace ] [Mesg: Released session REPOSITORY0]

Table 22 describes the types of IBM WebSphere ICS messages.

Table 22. Message types

Type	Description
Info	Informational only. You do not need to take action.
Warning	A default condition chosen by InterChange Server.
Error	A serious problem that you should investigate.
Fatal Error	An error that stops operation and should be reported.
Trace	Tracing information for the trace level specified.
Flow Trace	Flow tracing information for business objects.
Internal Error	A serious internal problem that should be investigated.
Internal Fatal Error	An internal error that stops operation. It should be reported.

**Note:** If a message with the Internal Error or Internal Fatal Error severity appears, record the circumstances surrounding the problem, and then contact IBM WebSphere ICS Technical Support.

## System logs

**InterChange Server logs:** InterChange Server can log messages to the following destinations:

- The system log file, `InterchangeSystem.log`, which you can create in the product directory.
- A log file that you specify.
- Standard output (STDOUT). Depending on your operating system, STDOUT appears in one the following ways:

### UNIX

STDOUT messages are redirected to the log file to `$PRODUCTDIR/logs/ics_manager.log`.

### Windows

STDOUT messages appear in the Command Prompt window in which InterChange Server starts.

- In addition to managing regular logging, InterChange Server can send an e-mail message to a specified user when it generates error or fatal error messages.

By default, trace messages are sent to the log file. In some cases, this file may become too large, so it is recommended that you specify a separate trace file. See “About tracing” on page 126 for instructions on setting up a separate file for trace messages.

“Configuring logging and tracing” on page 129 describes how to set the destination for logging.

To specify the recipient for e-mail notification, see “Configuring e-mail notification of log messages” on page 123.

“Log/Trace file management” describes how to keep log files from becoming too large.

**Connector Agent logs:** The connector agent and connector controller have separate mechanisms for logging. This section describes connector agent logging. Connector controller messages are sent to the log that contains the InterChange Server messages.

A connector agent logs messages to a local destination and can also send its messages to InterChange Server for logging. To specify a log file name, edit the `LogFile` property and insert the name of the log file you want to use. The default log file (located at `ProductDir/logs/connector_manager_ConnectoName.log` directory for UNIX and `STDOUT` for Windows), contains text for the error and informational messages raised from the connector. *Name* is the name of the application. For example, the default message file for the Oracle connector is `OracleConnector.txt`.

Table 23 describes the properties you can edit that determine where a connector agent logs messages.

Table 23. Connector agent log message properties

Property Name	Description	Type of Value
<code>LogAtInterchangeEnd</code>	Specifies whether the connector agent sends messages to InterChange Server in addition to logging them locally.  At InterChange Server, connector agent messages appear wherever server messages appear, according to the <code>InterchangeSystem.cfg</code> file.	Either true (sends messages to InterChange Server and enables e-mailing) or false (logs messages only locally).  The default value is false.
<code>LogFile</code>	Specifies where to write connector agent messages on the local system.	A file path or standard output (STDOUT). The default value is <code>STDOUT</code> .

For task instructions on configuring these properties, see “Configuring the connector agent logging destination” on page 133.

**High-Availability system logs:** The InterChange Server system log for High Availability (HA) is configured to reside on the cluster shared drive (for example, Z:). This log is viewable only from the active system. For information about setting up the log to reside on the shared drive, see the *System Installation Guide for UNIX or for Windows*.

The HA system also uses the Application log to provide information about the cluster and its events and failures. Be sure to check the Microsoft MSCS online help for information about this tool, and see the information in the section “Checking cluster status” on page 168.

## Log/Trace file management

When the InterChange system is started, a log file is created if one does not exist, or is appended to if it does. If the size of the log file is unlimited, it grows and its

size depends on the amount of time since it was last managed and the volume of transactions passing through the system. If a log file grows too large, you may not be able to open it or an application may require additional system resources to write to the files.

IBM WebSphere ICS system log files can be configured to a specified size and then automatically archived once they reach that size. As an added precaution, you can specify a number of archive files to use as a system backup. Each time the log file reaches its maximum size, the file is renamed as a new archive file. The archive file's name is derived from the original log or trace file name, with the following inserted into the name:

*\_Arc\_ number*

For example, using 5 archive files, if the log file has the name `InterchangeSystem.log`, the first archive created is named `InterchangeSystem_Arc_01.log`. When the new log file fills up, `InterchangeSystem_Arc_01.log` is renamed as `InterchangeSystem_Arc_02.log`, and the log file is again saved to `InterchangeSystem_Arc_01.log` and so on in a circular fashion, until there are five archive files. If there are five archive files, when a new log file is created, existing archive files are renamed and their numbers incremented so the number of archives matches the number you configured, then the oldest file, whose archive number is 05, is deleted. Figure 51 shows the progression of files using this configuration.

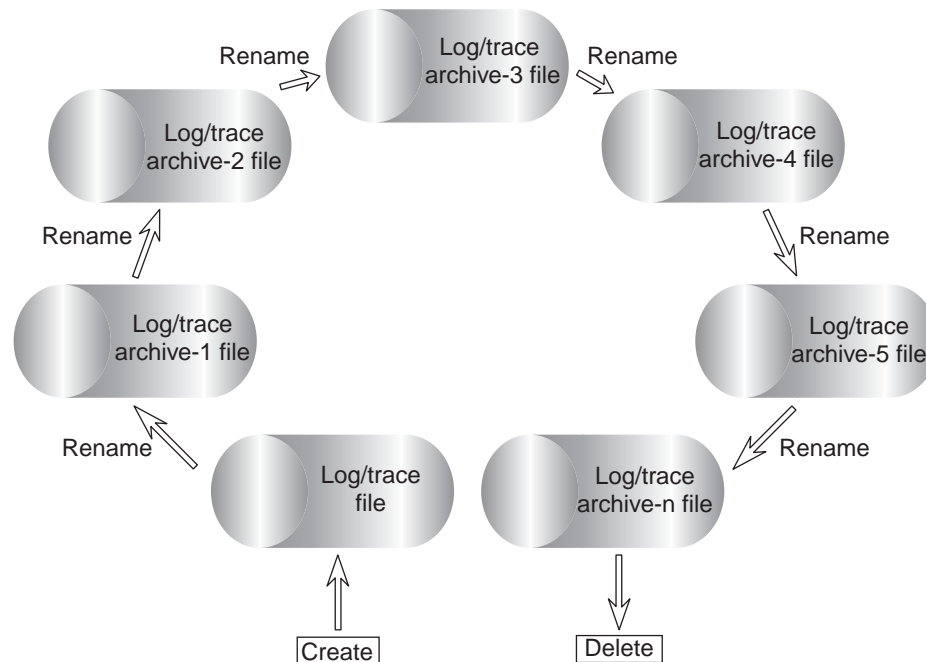


Figure 51. Circular archival logging

See the configuration tasks “Configuring logging and tracing” on page 129 for details.

If the system log file is configured for unlimited size, InterChange Server writes to the log until the disk that the log file is located on gradually fills, and if not administered to, finally produces an error message when the disk is full.

The data in these files should be deleted periodically:

- InterChange System log file. Contains a record of all interactions of the various system components with InterChange Server. The Edit Configuration tool allows you to specify the name and size of this file and a number of archive files for automatically saving versions of this file when it overruns its specified size. See “Configuring InterChange Server logging and tracing destinations” on page 130 for task information.
- Connector log file. Contains flows at the API level that were executed by the connector. Not all applications support these files.
- Connector error log file. Contains error messages returned by the connector. Not all applications support these files.

In addition to the above log files, other log files exist that are specific to each application. Most files are created during runtime if they don’t already exist. New information is appended to any existing file. Each component that supplies log information to the files must be taken down before proceeding with a backup.

Any file management procedure can be used, but the following periodic log file management is recommended:

- Rename the files by appending a date to the file.
- Move the files to an archive directory.

## Configuring e-mail notification of log messages

Error and fatal messages that are logged to the InterChange Server log can also be sent to the IBM WebSphere ICS system administrator, or any other recipient, by e-mail. By default, InterChange Server is configured to send e-mail notifications using JavaMail, but you can configure the server to send e-mail notifications using the e-Mail connector. For instructions on configuring e-mail notification using the e-Mail connector, refer to “Configuring e-mail notification at the system level.”

**Note:** If you want to configure collaborations for e-mail notification, you must use the e-Mail connector.

The following components can be configured for sending error and fatal messages to an e-mail recipient:

- InterChange Server (ICS)
- Collaboration objects
- Connectors

You can configure e-mail notification at the system level (set in the `InterchangeSystem.cfg` file), at the collaboration object level (set as a collaboration object property), or at the connector level (set as a connector property). If you configure e-mail notification at the system level, the configuration applies to all of the collaboration objects or connectors in the system. If you configure e-mail notification at the collaboration object or connector level, the configuration applies only to that specific component and supersedes the system configuration.

### Configuring e-mail notification at the system level

Configuring e-mail notification at the system level allows you to configure the server to send e-mail notifications for the following subsystems:

- Default e-mail recipient
- WebSphere MQ errors recipient
- Repository errors recipient

- Database persistence errors recipient
- Database access errors recipient
- Map errors recipient
- Security errors recipient
- Connector controller errors recipient
- DTP errors recipient
- Connector agent errors recipient

To configure e-mail notification at the system level, do the following:

1. From System Manager, right-click the server under Server Instances, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool in which you can edit the InterchangeSystem.cfg file.
2. Click the E-mail tab. A dialog box appears in the upper-right section of the System Manager window in which you can enter the parameters necessary for configuring e-mail notification at the system level (see Figure 52).

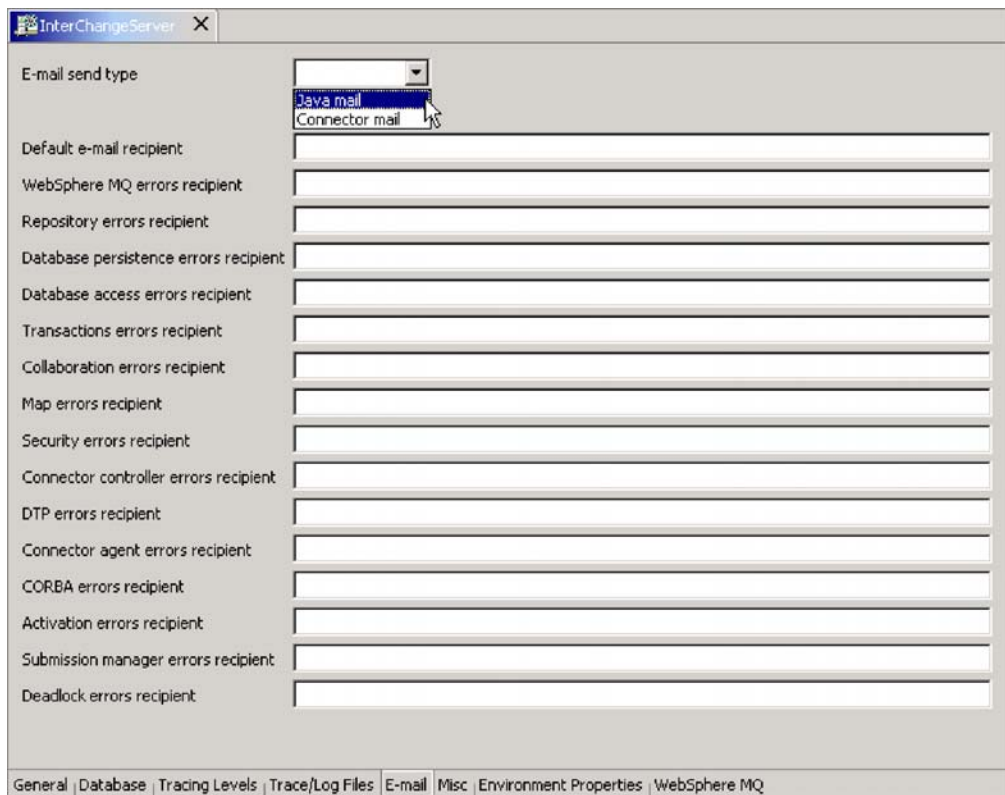


Figure 52. Edit Configuration, E-mail tab

3. Select an e-mail type from the "E-mail send type" drop-down menu. Your choices are "Java mail" or "Connector mail."
  - Select "Java mail" if you do not plan to configure collaborations for e-mail notification. If you plan to configure collaborations for e-mail notification, you must select "Connector mail." If you select "Java mail," an "SMTP mail host" field appears at the top of the list of subsystems. Type the host name of the SMTP mail server.

**Note:** Be sure to use the correct case, as this field is case sensitive.

- Select "Connector mail" if you plan to configure collaborations for e-mail notification. Selecting "Connector mail" requires that you install the e-Mail connector, then configure the e-Mail connector and e-Mail business object. For instructions on installing the e-Mail connector, see the *System Installation Guide for UNIX or for Windows*. For instructions on configuring the e-Mail connector and e-Mail business object, see "Configuring the e-Mail business object" and "Configuring the e-Mail connector."
4. Type one or more valid e-mail addresses in each subsystem field. The address must be SMTP-compliant. For information on SMTP compliancy, see "Using Valid E-mail addresses" on page 126.
  5. Select the "Save <server\_name>" option from the File drop-down menu of System Manager. The system-level e-mail notification information you entered is saved in the `InterchangeSystem.cfg` file.
    - If you chose "Java mail" as your e-mail send type, you can skip the following two sections: "Configuring the e-Mail business object" and "Configuring the e-Mail connector."
    - If you chose "Connector mail" as your e-mail send type, you must follow the instructions in the following sections: "Configuring the e-Mail business object" and "Configuring the e-Mail connector."

**Configuring the e-Mail business object:** The instructions in this section are required only if you are using the e-Mail connector. Configure the `EmailNotification` business object to hold the e-mail address of the person who receives e-mail if the intended e-mail recipient is unreachable. As a fail-safe, this should probably be the mail administrator, not the ICS system administrator, to ensure the mail is delivered if the ICS system administrator is unreachable.

1. From System Manager, right-click the `EmailNotification` business object, then select "Edit definition." The `EmailNotification Business Object Designer` window appears.
2. From the Attributes tab, type in the return e-mail address in the Defaults column of the `FromAddress` attribute.
3. Select Save from the File drop-down menu.
4. Close Business Object Designer.

**Configuring the e-Mail connector:** The instructions in this section are required only if you are using the e-Mail connector. To configure the e-Mail connector for sending e-mail:

1. From System Manager, double-click the `EmailConnector` object. The `Connector Configurator` window appears.
2. From the Connector Specific Properties tab, double-click the Value cell for the `SMTP_MailHost` property, then type in the name of the Simple Mail Transport Protocol (SMTP) host in the text field.
3. Click Save > To Project from the File drop-down menu.
4. Close the Connector Configurator window.

### **Configuring e-mail notification at the collaboration object level**

This section describes how to configure e-mail notification for a specific collaboration object. Configuring e-mail notification for collaborations requires that you use the e-Mail connector. Be sure to select "Connector mail" when configuring e-mail notification at the system level, and be sure to follow the instructions in the following sections to configure the e-Mail connector: "Configuring the e-Mail business object" and "Configuring the e-Mail connector."

**Note:** Configuration parameters set at the collaboration object level supersede those set at the system level.

To set an e-mail address to receive messages for a collaboration, do the following:

1. From System Manager, right-click the collaboration object for which you want to configure e-mail notification, then select Properties. The Properties dialog box appears.
2. In the Collaboration General Properties tab, enter a valid e-mail address in the "Email notification address" field. The address must be SMTP-compliant. For information on SMTP compliancy, see "Using Valid E-mail addresses."
3. Click OK to save your changes and close the window.
4. Restart the collaboration for changes to take effect.

### **Configuring e-mail notification at the connector level**

This section describes how to configure e-mail notification for a specific connector.

**Note:** Configuration parameters set at the collaboration object level supersede those set at the system level.

1. From System Manager, right-click the connector for which you want to configure e-mail notification, then select Edit Definition. The Connector Configurator window appears.
2. From the Properties tab, select true in the Value field of the LogAtInterchangeEnd property. This enables connector messages to be mailed to the InterChange Server log.
3. Click Save > To Project from the File drop-down menu.
4. Close the Connector Configurator window.
5. Restart the connector for the change to take effect.

### **Using Valid E-mail addresses**

E-mail notification in the WebSphere ICS system supports Simple Mail Transport Protocol (SMTP) mail messages, therefore, the e-mail recipient value in the InterchangeSystem.cfg file and the collaboration e-mail addresses must be standard Internet addresses.

A valid e-mail address entry can be one or more fully qualified Internet addresses, separated by semi-colons. For example, a valid entry for two recipients is:

```
JohnDoe@company.com;FredSmith@company.com
```

You cannot use personal address aliases, such as an alias defined in a personal address book. However, a valid address can be an alias defined in a mail server, such as Eng@company.com. In this case, the mail server decodes the alias and sends e-mail to all members of the alias. For example, a decoded alias might be person1@some\_company.com, person2@another\_company.com, and so forth.

## **About tracing**

To troubleshoot a problem, you can turn on tracing. Trace messages help you monitor actions taken in components of the IBM WebSphere ICS system. Trace levels define the amount of detail written to the trace file. The higher the trace level, the more detail you receive. Tracing differs from logging in the following ways:

- Logging always occurs, but tracing can be turned on and off as needed.
- Tracing contains more detailed information than logging about the state of components and the actions taken by them.



- Logging and tracing settings are persistent after reboots.

Tracing is off by default because it produces messages that are more detailed than you normally need. You can turn tracing on and off as necessary while InterChange Server is running.

This section covers the following topics:

“Tracing services of InterChange Server” on page 127

“Tracing collaboration objects” on page 127

“Tracing connectors” on page 128

“Tracing maps” on page 128

“Tracing business objects (flow tracing)” on page 128

“Tracing Web gateways” on page 129

“Tracing SNMP agents” on page 129

### Tracing services of InterChange Server

Tracing services for InterChange Server are initially set in parameters of the configuration file for InterChange Server (by default, this file is called `InterchangeSystem.cfg` and resides in the product’s top level directory. For details about these parameters, refer to “InterChange Server Configuration Parameters” in the *System Installation Guide for UNIX or for Windows*. Settings for these parameters can be updated in the Edit Configuration tool of System Manager, as described in “Configuring tracing levels for InterChange Server, business objects, and collaborations” on page 134.

### Tracing collaboration objects

You can trace the execution of a collaboration object. Tracing writes detailed messages about execution of the collaboration object to the log destination, which is specified in the `InterchangeSystem.cfg` file. Tracing collaborations is persistent. There are two trace level settings for collaborations, system level and collaboration level.

System level tracing returns runtime information for the collaboration. For example, if you want to trace the state changes of the collaboration, set the system trace level to 3.

You can set collaboration object tracing to one of the following levels:

---

1	Traces the receipt of business objects from connectors and the starting of the appropriate scenarios.
2	Prints messages for level 1. In addition, traces the start and completion of each scenario, reporting both forward execution and rollback.
3	Prints messages for levels 1 and 2. In addition, traces the execution of each scenario decision block or action.
4	Prints messages for levels 1 through 3. In addition, traces the sending and receipt of each business object by each scenario.
5	Prints messages for levels 1 through 4. In addition, traces the sending and receipt of each business object by each scenario, printing the value of each attribute in the business object.

---

For configuration task instructions, see “Configuring tracing levels for InterChange Server, business objects, and collaborations” on page 134.

## Tracing connectors

A connector contains two components, the connector controller and the connector agent. The two components can be in different locations on the network and are traced differently. For more information about tracing connectors, see the *Connector Development Guide for C++* and the *Connector Development Guide for Java*.

You can set connector agent and controller tracing to one of the following levels:

---

1	Traces initialization and the sending and receipt of business objects.
2	Prints messages for level 1. In addition, provides more details than Level 1 for the same types of events.
3	Prints messages for levels 1 and 2. In addition, traces the exchange of messages between the connector agent and the messaging driver.
4	Prints messages for levels 1 through 3. In addition, traces the passing of business objects between internal levels of the connector.
5	Prints messages for levels 1 through 4. In addition, traces the passing of administrative messages between internal levels of the connector.

---

A new or changed tracing level takes effect immediately.

For configuration task instructions, see “Configuring connector tracing” on page 135.

## Tracing maps

Tracing for IBM WebSphere ICS maps can be done through System Manager. Tracing maps is useful for debugging and keeping track of information as well as error messages created by the map. Tracing of IBM WebSphere ICS maps is turned off by default.

For more information about tracing maps, see the *Map Development Guide*. For configuration task instructions, see “Configuring map tracing” on page 136.

## Tracing business objects (flow tracing)

Business object trace logging provides a way to trace the progression of business objects from one processing point to another, based on notification messages that are generated at each point. For example, with level 2 tracing, when a business object arrives at a collaboration for processing, a trace message is logged.

Table 24 describes the configurable levels associated with business object tracing:

Table 24. Business object tracing levels

---

Level	Description
0	No tracing.
1	Event status (such as Successful or Failed) and event identity information.
2	Minimal event tracing. Information about when a business object enters/exits systems, such as connectors, maps, relationships, and collaborations. Includes level 1 information.
3	Provides event tracing from level 2 and a business object dump at entry/exit of systems. System performance impact.
4	Detailed tracing. Provides tracing for system components such as connectors, maps, relationships, and collaborations, as well as mapping these traces to level 3 event tracing. System performance impact.

---

For configuration task instructions, see “Configuring tracing levels for InterChange Server, business objects, and collaborations” on page 134.

## Tracing Web gateways

Web gateway tracing is provided at two levels, minimal and maximum. This tracing allows you to view information pertaining to whether communication processing is being performed correctly on the gateway. The gateway trace level is set from the Gateway Configuration Tool. For instructions, see the *System Implementation Guide*.

Table 25 describes the configurable levels associated with gateway tracing:

*Table 25. Gateway tracing levels*

Level	Description
0	No tracing.
1	Minimal tracing (such as bind requests, socket opens, and so forth).
5	Maximum tracing (such as HTTP/HTTPS requests, including all headers).

## Tracing SNMP agents

SNMP agent tracing provides trace information for checking the operation of the SNMP agent. The higher the level, the more verbose the output.

Table 26 describes the configurable levels associated with SNMP agent tracing:

*Table 26. SNMP agent tracing levels*

Level	Description
0	No tracing.
1	Message trace (displays Info, Warning, and Error type messages).
2	Low. Displays all incoming requests and outgoing messages.
3	Medium. Displays information such as Object ID, variable bindings and values.
4	High. Displays commands that access data on ICS and the values that are passed.
5	Method tracing. Errors resulting from ICS interface methods.

## Configuring logging and tracing

This section describes how to set up message logging and tracing. These settings can be made by using The Edit Configuration option in System Manager and by manually editing the `InterchangeSystem.cfg` file. The following tasks are described:

“Opening the Edit Configuration tool” on page 130

“Configuring InterChange Server logging and tracing destinations” on page 130

“Configuring other InterChange Server logging and tracing parameters” on page 132

“Configuring the connector agent logging destination” on page 133

“Configuring tracing levels for InterChange Server, business objects, and collaborations” on page 134

“Configuring the collaboration object trace level” on page 135

“Configuring connector tracing” on page 135

“Configuring map tracing” on page 136

## Opening the Edit Configuration tool

InterChange Server must be running and in the Connected state to use the Edit Configuration tool. The Edit Configuration tool can manage only one InterChange Server per session.

To open the Edit Configuration tool:

1. From System Manager right-click a server listed under Server Instances, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the InterchangeServer.cfg file.
2. Click either the Tracing Levels tab or the Trace/Log Files tab to configure tracing and message logging. See “Configuring InterChange Server logging and tracing destinations” and “Configuring tracing levels for InterChange Server, business objects, and collaborations” on page 134 for task instructions.

## Configuring InterChange Server logging and tracing destinations

Use the Edit Configuration tool to configure the destination for InterChange Server message logging and tracing.

When configuring these settings, keep the following information about STDOUT in mind:

### UNIX

If you set the logging and tracing to STDOUT, messages are automatically rerouted to `$PRODUCTDIR/logs/ics_manager.log`.

### Windows

If you plan to run InterChange Server as a service, you must set logging and tracing to file destinations. Setting logging and tracing to STDOUT prevents InterChange Server from being configured as a Windows service.

1. From the Edit Configuration tool, click the Trace/Log Files tab.  
The logging and tracing configuration window appears (see Figure 53 on page 131).

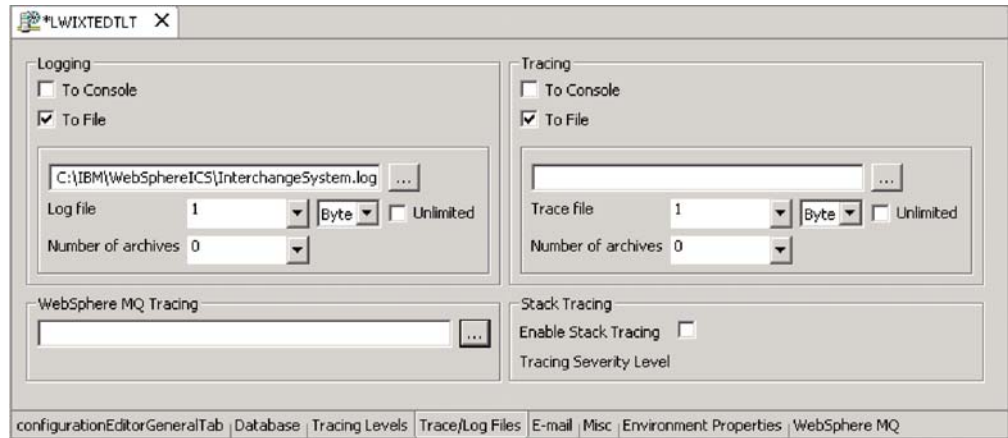


Figure 53. Edit Configuration tool, Trace/Log Files tab

2. In the Logging area, select the destination for system logging. You can log to both system console and a log file, but this option should be used only for debugging and on development systems.
  - a. To log to the system console (standard output), which is the default, make sure the To console (STDOUT) box is checked. To disable the console as the logging destination, uncheck the To console (STDOUT) option.
  - b. To log to a file, click the To File box and either type in the full pathname of the file or click the browse (...) button to navigate to the log file.
  - c. If you select a log file, optionally configure the size of the file in MBs (or keep the default, Unlimited) and the number of archive files to create. For information about archives, see “Maintaining event archives” on page 145.
3. In the Tracing area, select the destination for system tracing.
  - a. To log to the system console (standard output), click the To console (STDOUT) box. To disable the console as the logging destination, uncheck the To console (STDOUT) option.

If you choose to log messages to STDOUT, the messages appear in one of the following ways:

#### UNIX

If you set the logging and tracing to STDOUT, messages are automatically rerouted to `$PRODUCTDIR/logs/ics_manager.log`.

#### Windows

STDOUT appears in the Command Prompt window in which InterChange Server starts.

- b. To log to a file, click the To File box and either type in the full pathname of the file or click the browse (...) button to navigate to the trace file.
  - c. If you select a trace file, configure the size of the file in MBs (or keep the default, unlimited checked) and the number of archive files to create. For information about archives, see “Maintaining event archives” on page 145.
4. To configure WebSphere MQ tracing, accept the default pathname of the file or click the browse (...) button to navigate to the log file.
5. Click OK to save changes and exit.

Your changes take effect immediately, and if you already had a log file configured, it is saved and dated, and a new file created using the newly entered configuration.

## Configuring stack trace

When ICS stack tracing is enabled, the stack trace information is printed to the ICS log file. This can be useful for troubleshooting a newly installed system.

The following instructions describe how to configure the stack tracing feature.

1. From the InterChange Server Component Management view in System Manager, right-click the ICS instance for which you want to configure stack tracing, then select Edit Configuration.

The upper-right quadrant of the System Manager becomes a tool in which you can edit the `InterchangeSystem.cfg` file (see Figure 54).

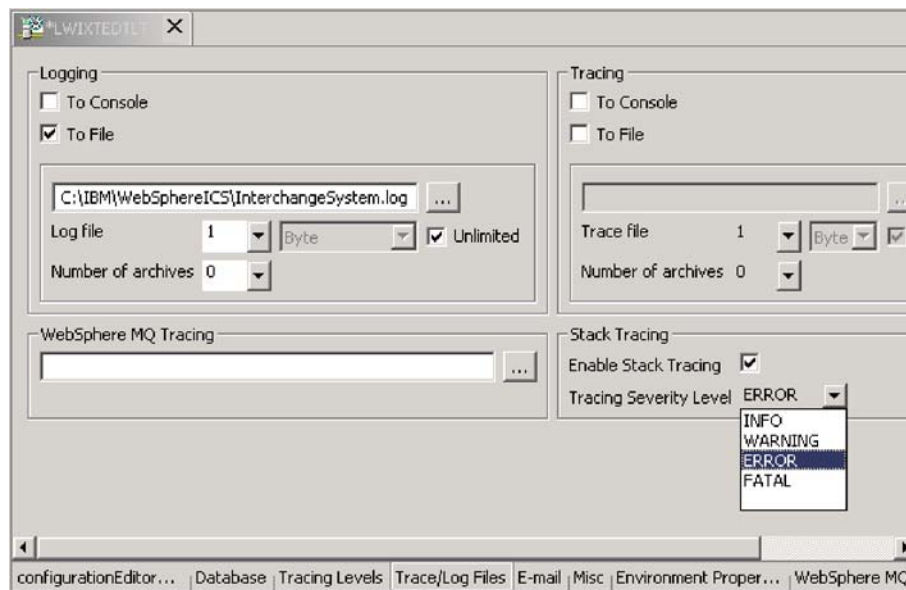


Figure 54. Configuring stack tracing

2. In the Trace/Log Files tab, enter the following values:
  - Select the Enable Stack Tracing box. When you enable stack tracing, the Tracing Severity Level option becomes visible.
  - In the Tracing Severity Level drop-down menu, select the severity level for which you want to print the stack tracing. The choices are: INFO, WARNING, ERROR, and FATAL. When stack tracing is printed to the log file, it will print the level you select and all levels above it. For example, if you select ERROR, both ERROR and FATAL messages will be printed to the log file.

## Configuring other InterChange Server logging and tracing parameters

You can set other logging and tracing parameters, such as file size and number of archived files. The following instructions describe how to perform these tasks.

To set the maximum size of the log and trace files:

1. From the InterChange Server Component Management view in System Manager, right-click the ICS instance for which you want to set the maximum size for log and trace files, then select Edit Configuration.

The upper-right quadrant of the System Manager becomes a tool in which you can edit the `InterchangeSystem.cfg` file (see Figure 54 on page 132).

2. To limit the log file, in the Trace/Log Files tab, select a number to represent the maximum log file size from the Log file drop-down menu, then select Byte, KB (kilobyte), MB (megabyte), or GB (gigabyte).

**Note:** The To File and the Unlimited checkboxes in the Logging section must both be selected before you can enter a maximum log file size.

3. To limit the trace file, in the Trace/Log Files tab, select a number to represent the maximum trace file size from the Trace file drop-down menu, then select Byte, KB (kilobyte), MB (megabyte), or GB (gigabyte).

**Note:** The To File and the Unlimited checkboxes in the Tracing section must both be selected before you can enter a maximum trace file size.

4. Stop and restart InterChange Server.

To configure the number of archive files:

1. From the InterChange Server Component Management view in System Manager, right-click the ICS instance for which you want to set the maximum size for log and trace files, then select Edit Configuration.

The upper-right quadrant of the System Manager becomes a tool in which you can edit the `InterchangeSystem.cfg` file (see Figure 54 on page 132).

2. To set the number of log or trace file archives, from the Trace/Log Files tab, select a number from the "Number of archives" drop-down list from either the Logging or Tracing section.

**Note:** The To File and the Unlimited checkboxes must both be selected before you can select a number from the "Number of archives" drop-down menu. This is true for either the Logging or Tracing section.

3. Stop and restart InterChange Server.

### Configuring the connector agent logging destination

The two connector components have separate mechanisms for logging. Connector controller logging is sent to the `InterchangeServer.log` file. This section describes configuring the connector agent log file name and location.

For background information on connector agent logging, see "Connector Agent logs" on page 121.

You can set the destination for connector agent logging using one of the following methods:

- If you are installing the Remote Agent, you can configure the logging destination during the installation. This option is available for both UNIX and Windows operating systems. For more information about installing Remote Agent, see the *System Implementation Guide*.
- At any other time, you can use Connector Configurator to configure the connector agent logging destination. This can be done only on a Windows 2000 machine.

To configure the destination for connector agent logging using Connector Configurator, connect to a server, then follow these steps:

1. From System Manager, right-click a connector object, then select Edit Definitions. The Connector Configurator window appears.

- From the Standard Properties tab, select one of the logging property values (see Table 23 on page 121 for an explanation of these values), then click Edit.
- Enter the new value in Value field, then click OK.

For example, change the LogAtInterchangeEnd value to true to send messages to the InterChange Server log. If InterChange Server is configured to send e-mail when error and fatal messages are logged, e-mail is sent for the connector agent messages as well.

- Repeat Steps 3 and 4 to edit the other logging property if necessary.  
For example, to send connector agent messages to a message file instead of the default STDOUT, enter the full pathname of the file in the Value field.
- Restart the connector for changes to take effect.

### Configuring tracing levels for InterChange Server, business objects, and collaborations

This section describes how to configure tracing levels for business objects, collaborations, and the IBM WebSphere ICS subsystems using the Edit Configuration tool in System Manager. For details on viewing the trace messages, see “Log Viewer and tracing” on page 137.

To open the Edit Configuration tool, right-click the server from the Server Instances section of System Manager, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the InterchangeServer.cfg file.

To configure tracing:

- From Edit Configuration tool, select the Tracing Levels tab. Two categories appear in the window: Flow Tracing Levels and IBM WebSphere Business Integration System Tracing Levels (see Figure 55).

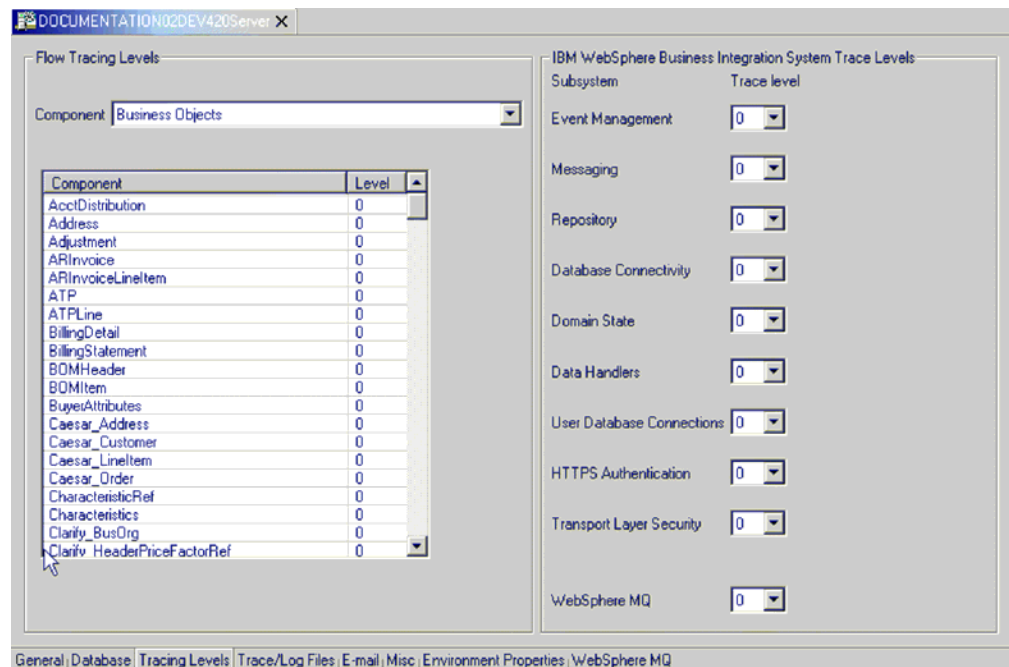


Figure 55. Edit Configuration tool, Tracing Levels tab



2. To configure subsystem tracing, in the IBM WebSphere Business Integration Trace Levels section, click the down arrow next to the subsystems you want to trace and set the trace level using the down arrow menu.

Setting these trace levels updates parameters in the `InterchangeSystem.cfg` file. For details about what information is produced at the various tracing levels, see “InterChange Server Configuration Parameters” in the *System Installation Guide for UNIX or for Windows*.

3. To configure tracing for collaborations, in the Flow Tracing Levels section, click the Component down arrow and select Collaborations.

The collaboration names that are configured for the system display in the list box.

Click the down arrow in the Level area to select the trace level for that collaboration.

For a description of the trace levels for collaborations, see “Tracing collaboration objects” on page 127.

4. To configure flow tracing for business objects, click the Component down arrow and select Business Objects. For a description of tracing levels for business objects, see “Tracing business objects (flow tracing)” on page 128.  
The business object names that are configured for the system display in the list box.
  - a. Click the down arrow in the Level area to select the trace level for that business object.
5. Click OK to save changes and exit.  
The trace levels for the subsystems, business objects and collaborations are immediately in effect.

### **Configuring the collaboration object trace level**

This section describes one of two methods that can be used to configure collaboration object trace levels. For the alternative method, see “Configuring tracing levels for InterChange Server, business objects, and collaborations” on page 134.

To configure the runtime system trace level for a collaboration object:

1. From System Manager, right-click the collaboration object name, then select Properties. The Properties dialog box appears.
2. Select a System trace level value, then click OK.

Collaboration object tracing returns messages from inside the collaboration. For a description of collaboration object trace levels, see “Tracing collaboration objects” on page 127.

To configure collaboration object tracing:

1. From System Manager, right-click the collaboration object name, then select Properties. The Properties dialog box appears.
2. In the Collaboration Trace Level field, select a value, then click OK.

A collaboration object starts tracing as soon as its tracing level changes.

### **Configuring connector tracing**

For background information, see “Tracing connectors” on page 128.

**Connector controller tracing:** To trace a connector controller, set the `ControllerTraceLevel` property value to the trace level. Any changes to this

property take effect immediately. Trace messages for connector controllers appear wherever InterChange Server sends its trace messages.

To set a connector controller's tracing level, do the following:

1. From System Manager, right-click the connector object, then select Edit Definitions. The Connector Configurator window appears.
2. From the Standard Properties tab, click in the Value field of the ControllerTraceLevel property, then set the controller tracing level.
3. Select Save > To Project from the File drop-down menu.
4. Close Connector Configurator.

**Connector agent tracing:** To trace a connector agent, set the AgentTraceLevel property value to the trace level. Any changes to this property take effect immediately. Trace messages for connector agents appear wherever the connector agent logs messages. To set a connector agent's tracing level, do the following:

1. From System Manager, right-click the connector, then select Edit Definitions. The Connector Configurator window appears
2. From the Standard Properties tab, click the Value field of the AgentTraceLevel property, then set the agent tracing level.
3. Select Save > To Project from the File drop-down menu.
4. Close Connector Configurator.

## Configuring map tracing

For background information, see "Tracing maps" on page 128. To set the trace level for a map:

1. From System Manager, right-click the map object, then select Properties. The Maps Property Page appears (see Figure 56).

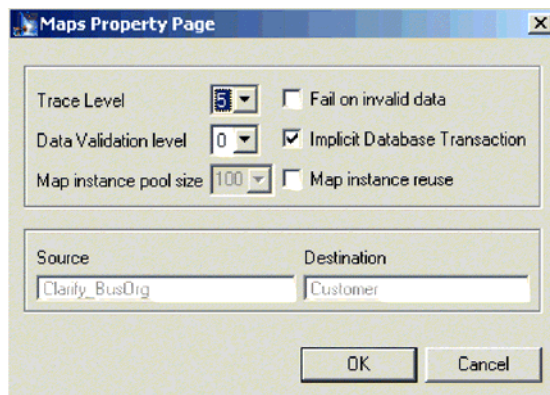


Figure 56. Maps Property Page

2. In the Trace Level field, enter the appropriate value.
3. Click OK.

## Using tracing

To troubleshoot a problem, you can turn on tracing. Trace messages help you monitor actions taken in components of the WebSphere ICS system. Trace levels define the amount of detail written to the trace file. The higher the trace level, the more detail you receive.

Tracing is off by default because it produces messages that are more detailed than you normally need. You can turn tracing on and off as necessary while InterChange Server is running.

For background information about tracing, see “About tracing” on page 126. For configuration information about tracing, see “Configuring logging and tracing” on page 129.

### Log Viewer and tracing

Use Log Viewer to display trace information for InterChange Server. In addition to displaying the debugging trace information for collaborations, it allows you to view the progression of a business object as it passes from one processing point to another, for example as the business object exits the collaboration processing point and is sent on to other collaborations or connectors for processing or is forwarded to the mapping stage for data transformation. For information about flow tracing, see “Tracing business objects (flow tracing)” on page 128.

**Note:** Log Viewer is a WebSphere ICS tool, which means it runs only on Windows 2000 or Windows XP machines. To configure or view a UNIX log file using Log Viewer, copy the log file from the UNIX machine to a Windows machine that has the WebSphere ICS product installed.

## Viewing log messages

InterChange Server system message logging is used to communicate messages, component state changes, and failures.

**Note:** Log Viewer views log files and therefore does not need InterChange Server to be running. The WebSphere ICS system administrator must have the appropriate file system permissions set to view log files.

You can view log files containing messages and explanations of system messages either of these ways:

- Use Log Viewer to view system log files in its window, as well as obtain a message explanation interactively by clicking on a message number. See “Using Log Viewer.”
- Use a text editor to view a log file and to search the IBM WebSphere ICS message file by message number or text content to obtain an explanation.

### Using Log Viewer

Log Viewer allows you to see all messages contained in a log file. You can sort and filter the output display as well as print, save, and e-mail of the file.

**Note:** Log Viewer is a WebSphere ICS tool, which means it runs only on Windows 2000 or Windows XP operating systems. To configure or view a UNIX log file using Log Viewer, copy the log file from the UNIX machine to a Windows machine that has the WebSphere ICS product installed.

To start Log Viewer, you can either:

- Select Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration Toolset > Administrative > Log Viewer. Use the Open option of the File menu to browse for the log file.
- Use the Run command from the Start menu and browse for the LogViewer.exe file. Use the Open option of the File menu to browse for the log file.

Using the Log Viewer menu options, you can perform the following tasks:

- “Setting Log Viewer Preferences”
- “Changing how messages are viewed” on page 140
- “Manipulating the Log Viewer display output” on page 141
- “Flow tracing a business object” on page 142

Log Viewer, displaying a sample log file is shown in Figure 57.

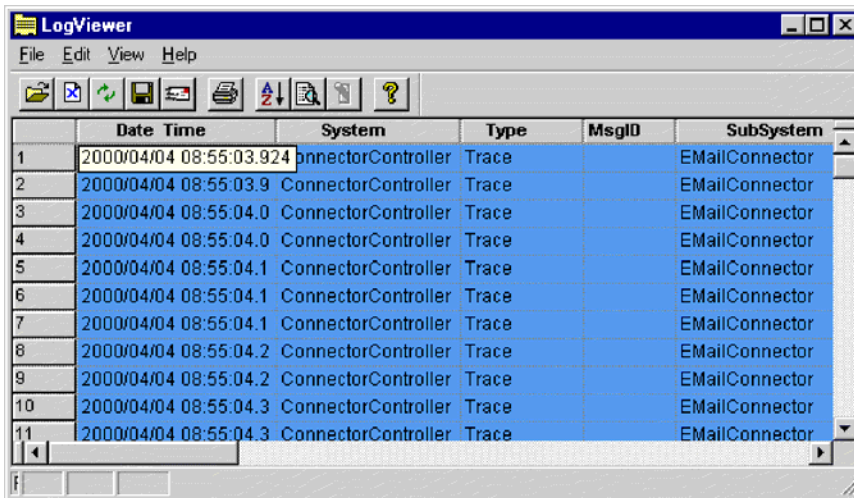


Figure 57. Log Viewer

#### Setting Log Viewer Preferences:

1. Log Viewer preferences are set by selecting the Preferences option from Log Viewer Edit menu.

The User Configuration Options, General properties window appears (see Figure 58).

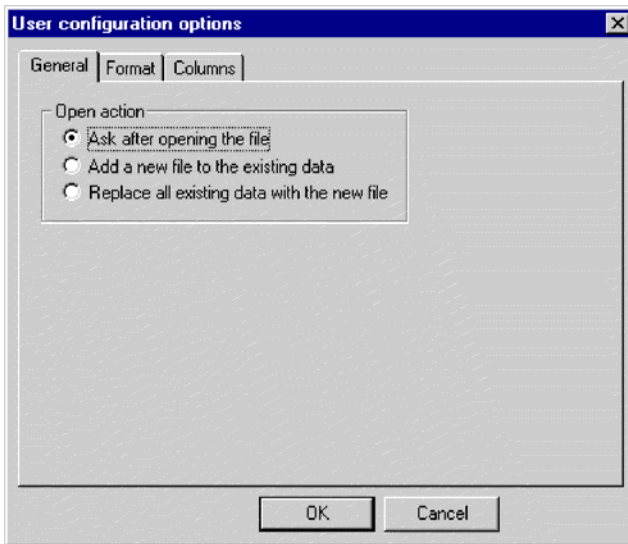


Figure 58. Log Viewer User configuration options dialog box, General Properties tab

This window allows you to determine how to display the log file when you open a log file. The available choices are:

- Query about what to do each time you open a log file.

- Merge the log file you are opening with the log file that is currently displayed.
  - Replace the log file that is currently displayed with the contents of the one you are opening.
2. To change the background color and font of the Log Viewer messages, click the Format tab.

The User Configuration Options, Format properties window displays (see Figure 59).

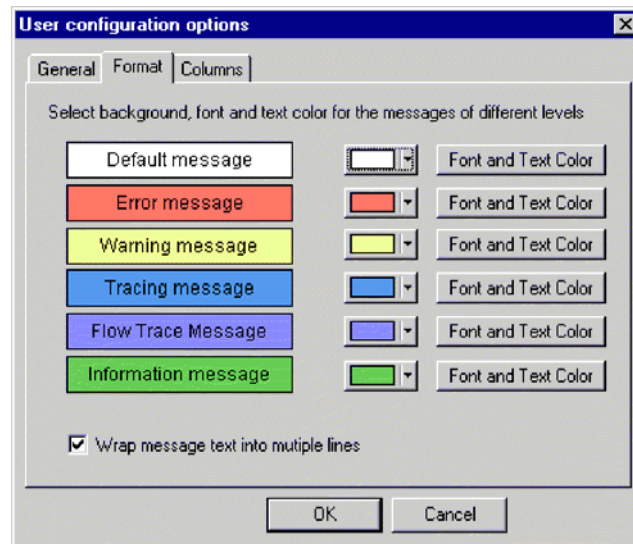


Figure 59. Log Viewer User configuration options, Format tab

This window allows you to determine how to display the log messages. The available choices are:

- Assign different background colors and fonts for each of the types of messages that display so you can easily recognize their severity (for example, red background with larger font allows for Warning messages).
  - Wrap the text of messages if the text is wider than the column.
3. To change which Log Viewer columns display, click the Columns tab.
- The User Configuration Options, Columns properties window displays (see Figure 60 on page 140).

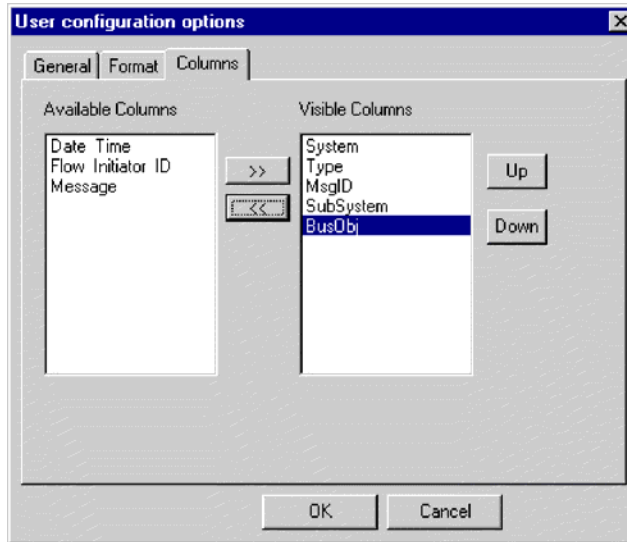


Figure 60. Log Viewer User configuration options, Columns tab

This window allows you to determine which columns display in Log Viewer. To change the columns that display:

- To display a column, highlight a column name in the Available Columns pane and click the >> button to move it to the Visible Columns pane.
- To hide a column, highlight a column name in the Visible Columns pane and click the << button to move it to the Available Columns pane.
- Click any of the column names in the Available Columns pane and click the Up or Down button to change its ordering from left to right in the Log Viewer display. Up moves columns to the left and Down moves columns to the right.

**Changing how messages are viewed:** The View menu contains additional options to change Log Viewer displays. In that menu, you'll find the following options:

- Display/hide the Log Viewer toolbar.
- Display/hide the Log Viewer status bar.
- Split the window into two or more views
- Filter or show all messages. Configure filter options by checking filtering options in the filter tabs, such as time range or by type of message (see Figure 61 on page 141 and Table 22 on page 120). In the Activate Filters area, click the box that is associated with the tab where you selected the filter options, and click OK to enable filtering. The filtered output can be toggled on or off with the Filter Toggle button on the Log Viewer toolbar.

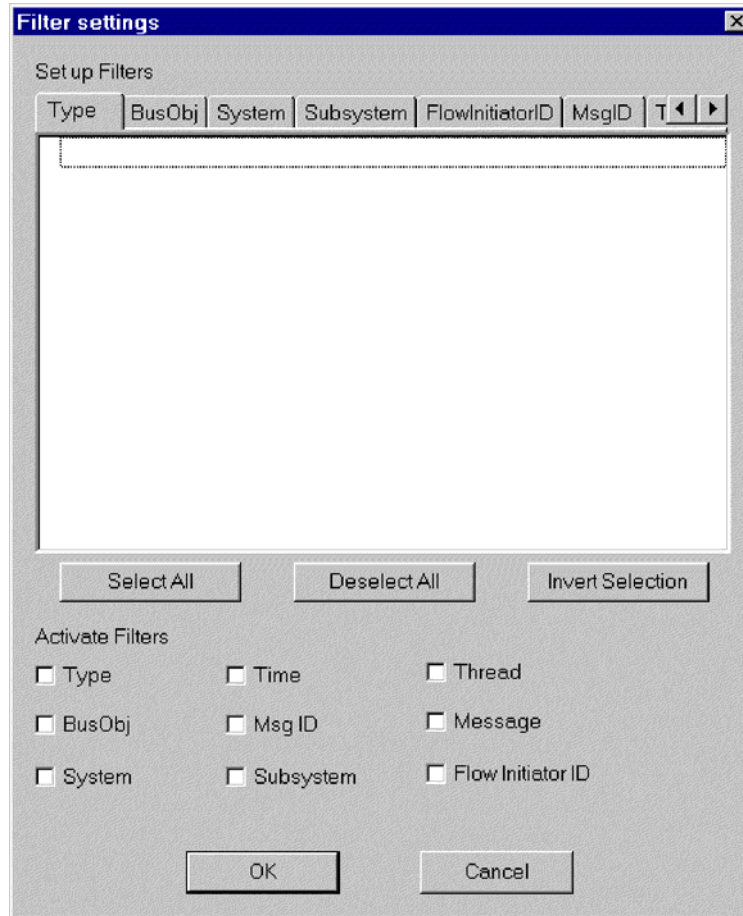


Figure 61. Log Viewer Filter settings Window

- Sort the messages; Figure 62 shows the Sort options.  
Click the down arrow in each sort field to select Date/Time, EventID, or business object. Further sort by ascending or descending ordering.

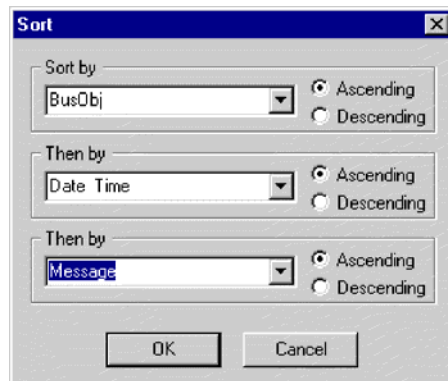


Figure 62. Log Viewer Sort Window

**Manipulating the Log Viewer display output:** Several options are available for manipulating Log Viewer output. In the File menu, there are options for print previewing, printing, saving, refreshing the display, sending to an e-mail recipient, and determining the style for page setup, headers and footers. The variables for header and footers are:

\$F	Name of file
\$A	Application name
\$P	Page number
\$N	Total number of pages
\$D	Date (can be followed by additional parameters (for example \$D{%y:%h:%m}))

**Filtering messages:** To filter the messages that will be displayed in Log Viewer, choose View->Filter->Use Filter in the Log Viewer menu. The Filter Settings dialog will display.

The Filter Settings dialog displays categories that correspond to the parameters of the logging message format (see Table 22 on page 120 for a descriptive list of the parameters).

In the Filter Settings dialog, you first choose the filtering categories that you want to use, then select the specific items that you want to display from each category, and then choose which filters you want to activate for your current Log Viewer display.

Follow these steps:

1. In the Filter Settings dialog, choose a tab under Set up Filters to display the items that you want to use for filtering messages. For example, choose the BusObj tab to display a list of business objects to be used in filtering, and choose Time if you want to filter according to the timestamp of the message. You can set up multiple filters, and use them either separately or in conjunction with each other.
2. In the displayed list of items, select each item for which you want to view messages in Log Viewer. For example, if you want to view messages related only to the Cost and Customer business objects, select only those business objects in the list. If you want to view only messages that are timestamped between 5 March 2002 at 9:00 AM and 6 March 2002 at 5:00 PM, select the range for those times under the Time tab.  
You can use the buttons below the list box to select all the displayed items, or to deselect all the displayed items, or to invert your current selection choices.
3. Under Activate Filters, check the box for each filter type that you want to activate. For example, if you want to see all messages for the Cost and Customer business objects (which you specified in the previous step), activate only the BusObj filter. If you want to see only those messages for the Cost and Customer business objects that have a particular timestamp, activate both the BusObj filter and the Time filter.
4. Choose OK. The Filter Settings dialog closes, and the Log Viewer display refreshes to show only those messages that you have allowed through the filters.

Note that in addition to filtering according to the categories, you can also display only those messages that contain a specific text string. To do so, choose Message under Set up Filters, enter the specific text for which you want to show messages, and check the box for Message under Activate Filters.

**Flow tracing a business object:** Flow tracing a business object or access flow allows you to track its progress throughout each of the processing points in its life cycle. Using Log Viewer, you can follow the progress by checking the trace



messages that display. Each business object has an flow initiator ID associated with it for just this purpose. If you sort the Log Viewer display by flow initiator ID and date/time, the trace messages for the business object are grouped together so you can easily follow its status. Sort by ascending or descending order to see a historical perspective or the latest event displayed first.

**Note:** Flow tracing is performed only while the business object is within the domain of InterChange Server, that is, from the connector controller of the source application to the connector controller of the destination application. Business object flow tracing is not performed while the business object is being processed by connector agents or applications.

To trace a business object flow:

1. Set the trace log file destination, if necessary (see “Configuring logging and tracing” on page 129).
2. Select the originating triggering business object (not the generic business object) to trace and set its trace level (see “Configuring tracing levels for InterChange Server, business objects, and collaborations” on page 134).
3. Send some event from the source connector to the destination connector.
4. Open Log Viewer (see “Using Log Viewer” on page 137).
5. Set the display preferences to view the flow tracing (see “Changing how messages are viewed” on page 140).
6. Click any error message button in the MsgID column to view the text of the message.
7. Click any of the business object name buttons in the BusObj column to view the data contained within the business object.

This action uses the Business Object Viewer, which allows you to save the data to a separate file. The file can then be read by either the Mapping tool or the Test Connector.

---

## Controlling server memory usage

OutOfMemoryExceptions thrown by the JVM can halt ICS. To reduce such occurrences, a “memory checker” feature in ICS can be used to pause the connectors when a predetermined level of memory usage is reached, ensuring that no new events are delivered to the connector, but still allowing it to process pending service call requests from ICS. This can reduce memory usage. The parameters for the memory checking feature can also control the speed of event delivery to InterChange Server from all connectors, providing finer granularity of flow control.

**Note:** The memory checker feature can also *negatively* affect the performance of ICS server, and should be used with caution. You should first try other methods of flow control, and use the memory checker only as the last solution for memory problems.

To implement memory checking, add parameters to the ICS startup script file (.bat file or .sh file) and to the ICS configuration file:

- Add the following parameter to the ICS startup script file, as a Java property parameter (-D option):  
CW\_MEMORY\_MAX

This value determines the maximum heap memory ICS is allowed to use. It should be the same value, in megabytes, as the JVM maximum heap size specified in the startup scripts of ICS (-Xmx parameter). For example:

`-DCW_MEMORY_MAX=512m`

- Add the following parameters to the [SERVER\_MEMORY] subsection of the ICS configuration file:

Parameter	Description
MEMORY_CHECK_SLEEP	This parameter controls the time interval (in minutes) that the memory checker thread sleeps. That is, it determines the time interval between two memory check operations. When this value set to zero (the default value), the memory checker does not work (meaning this flow control function is disabled). Example: <code>MEMORY_CHECK_SLEEP = 1</code>
MEMORY_UPPER_THRESHOLD_PCT	This parameter tells the memory control subsystem that when the current heap usage reaches this percentage of <code>CW_MEMORY_MAX</code> , the connectors should be paused. For instance, if the value for this parameter is set to 90 (the default value), and <code>CW_MEMORY_MAX</code> is set to 512m, when heap usage reaches 460 MB, all connectors will be paused.
MEMORY_LOWER_THRESHOLD_PCT	This parameter controls the speed of event delivery from connectors to InterChange Server. Its value is a percentage of <code>CW_MEMORY_MAX</code> (the default value is 80). When the current heap usage crosses this lower end threshold, the connector controller's event listener sleeps for a period of time. (The maximum amount of sleep time for the connector listener is set in the <code>SLEEPTIME_AFTER_MAX_THRESHOLD</code> parameter.) The sleep time varies depending how far the heap usage has exceeded the lower end threshold (but is still below the high threshold). This controls the speed at which events are delivered to ICS, providing a flow control mechanism. Example: <code>MEMORY_LOWER_THRESHOLD_PCT = 75</code>
SLEEPTIME_AFTER_MAX_THRESHOLD	This parameter controls the minimum time interval (in minutes) for the connector's pause time. The default value is 5. When the current heap usage reaches the predefined higher end of memory threshold, the connector sleeps for at least this period of time. After this period, the memory checker thread checks the current heap usage again, and, depending on whether or not the memory is over the threshold, pauses or resumes the connectors. This parameter also controls the maximum sleep time of the connector's event listener. Example: <code>SLEEPTIME_AFTER_MAX_THRESHOLD = 2</code>

---

## Maintaining event archives

Some connectors use event archives at the application. If a connector archives events, you can use the archive to troubleshoot problems with the connector. The application administrator probably needs to flush the archive periodically so that it does not grow too large.

An event archive can provide a useful backup role. If a collaboration stops running, its subscriptions are cancelled. While the collaboration is inactive, the connector continues to retrieve the associated events, but finds no active subscriptions, so it archives the events. When the collaboration restarts, you can examine the event archive and move archived events back to the event notification mechanism for reprocessing.

If a particular connector archives events, information on administration issues is available in that connector's documentation.

**Attention:** If you use WebSphere MQ for messaging, do not back up the MQ queues when performing regular backup procedures. The data in these queues in dynamic and, therefore, represents only in-progress transactions.

---

## Managing WIP connections

You should configure enough database Work-in-Progress (WIP) connections to accommodate potential runtime-driven work-load.

To do so, add the following property to the [Event\_Management] section of the InterchangeSystem.cfg file:

```
MIN_CONNECTIONS = n
```

Where

```
n=sum_of_all_max_threads_for_collabs_used_simultaneously_plus_num  
_connectors
```

The `sum_of_all_max_threads_for_collabs_used_simultaneously` is the sum of the Maximum number of concurrent events settings of all the collaboration objects.

Setting a sufficient minimum number of WIP connections helps eliminate the possibility of a server crash in circumstances where the server is attempting to reboot, and all connectors are trying to do WIP work while all collaboration threads are undergoing wip-commit.

---

## Managing failures

Managing failures in the IBM WebSphere ICS system consists of using troubleshooting resources to resolve problems. Critical errors that may cause events to fail can occur with a system component such as a connector or collaboration, or a third party component such as an integrated application.

When an error causes an event to fail to process properly, the WebSphere ICS system has built-in capabilities that let you resolve the problems. The WebSphere ICS system can be set up to pause collaborations within InterChange Server if a failure occurs. This section covers the following topics:

“Failure recovery for service calls” on page 146

“Strategies for InterChange Server recovery” on page 146

“Critical errors” on page 148

“Lost connection to application” on page 150

“Unknown connector agent status” on page 150

“Database connection failures” on page 151

“Flow failures” on page 151

## Failure recovery for service calls

To avoid sending duplicate events to destination application, you may want to prevent a recovery from automatically resubmitting all service calls that were in transit when a failure occurred. To do so, prior to the server failure, you can configure a nontransactional collaboration to persist any service call event in the In-Transit state when a failure and recovery occurs. When InterChange Server recovers, the service call events remain in the In-Transit state, and you can use the Unresolved Flows dialog to examine individual failed events and control when (or if) they are resubmitted.

To configure a collaboration to persist a failed service call in-Transit state, set the Persist Service Call In Transit State checkbox of the Collaboration General Properties window.

**Attention:** If the collaboration is part of a collaboration group, setting this property will implicitly apply the same property setting to all other collaborations in the group. You should not set one collaboration for Persist Service Call In Transit State unless you want all other collaborations in that collaboration group to have that setting as well.

**Attention:** For transactional collaborations, it is recommended that you not set this property; leave the Persist Call In Transit State checkbox blank.

## Strategies for InterChange Server recovery

If InterChange Server (ICS) fails while processing events, all the events currently in the Work-in-Progress (WIP) queue need to be either recovered or otherwise dealt with when the server reboots. Potentially, because of memory requirements, the recovery of the WIP events can slow or even halt the server reboot. The IBM WebSphere ICS product provides two features—deferred recovery and asynchronous recovery—for improving the time it takes the server to reboot and for making the server available for other work before all events have been recovered.

In release 4.2 of the product, two new features were added that can assist in the efficiency of deferred recovery and asynchronous recovery: flow control and storing business object keys as part of the WIP data. Both features reduce the amount of memory needed during an ICS recovery, and therefore, should decrease significantly the amount of time necessary for ICS to reboot during a recovery.

Storing business object keys as part of the WIP data means that during recovery, the business object key is retrieved without deserializing the business object, thus preventing the need to do an MQ or a database round-trip. Flow control is a service that allows you to configure either system-wide or component-level queue

depth parameters in an effort to control the memory demands on ICS. For more information about configuring flow control, see “Configuring system-wide flow control” on page 54.

This section covers the following topics:

“Deferred recovery of collaboration events” on page 147

“Asynchronous recovery” on page 148

### **Deferred recovery of collaboration events**

In deferred recovery, recovery of a collaboration’s WIP events is deferred until after the server has rebooted, thereby saving the memory usage associated with those events.

After the server has rebooted, you can resubmit the events manually. Note the following recommendations:

- To avoid the possibility of data corruption due to event sequencing, resubmit the deferred events before you process any new events for that collaboration.
- Resubmit the deferred events in their original sequence.

You establish deferred recovery by setting the RECOVERY\_MODE property of a collaboration object.

**Attention:** Use of deferred recovery compromises the temporal sequencing of events, which may cause data integrity problems. Use this feature only if sequencing is not important to you.

The RECOVERY\_MODE property has two settings, which have these behaviors when a server failure and reboot occurs:

- **Deferred**  
Events that were in the Working state before the server failure will be transitioned to Deferred. No events for this collaboration will be recovered until you resubmit them manually.
- **Always**  
Events that were in the Working state before the server failure will be recovered. If there are any existing events that were previously transitioned to the deferred state, they will remain deferred until you resubmit them.

The default setting is Always.

**Note:** Changing the collaboration recovery mode value from Deferred to Always does not cause existing deferred events to be recovered, nor does it change existing deferred events to the Working state. Events in the deferred recovery state will remain deferred until and unless you manually resubmit them.

To set the Collaboration Recovery mode value when you create a new collaboration object, or to change the Recovery mode value at a later time, you can use the Collaboration General Properties window while InterChange Server is running. To open the window and change the value, do the following:

1. From System Manager, right-click a collaboration object, then choose Properties. The Properties dialog box appears (see Figure 63 on page 148).

- In the Collaboration Properties dialog, choose the Collaboration General Properties tab. The following dialog displays:

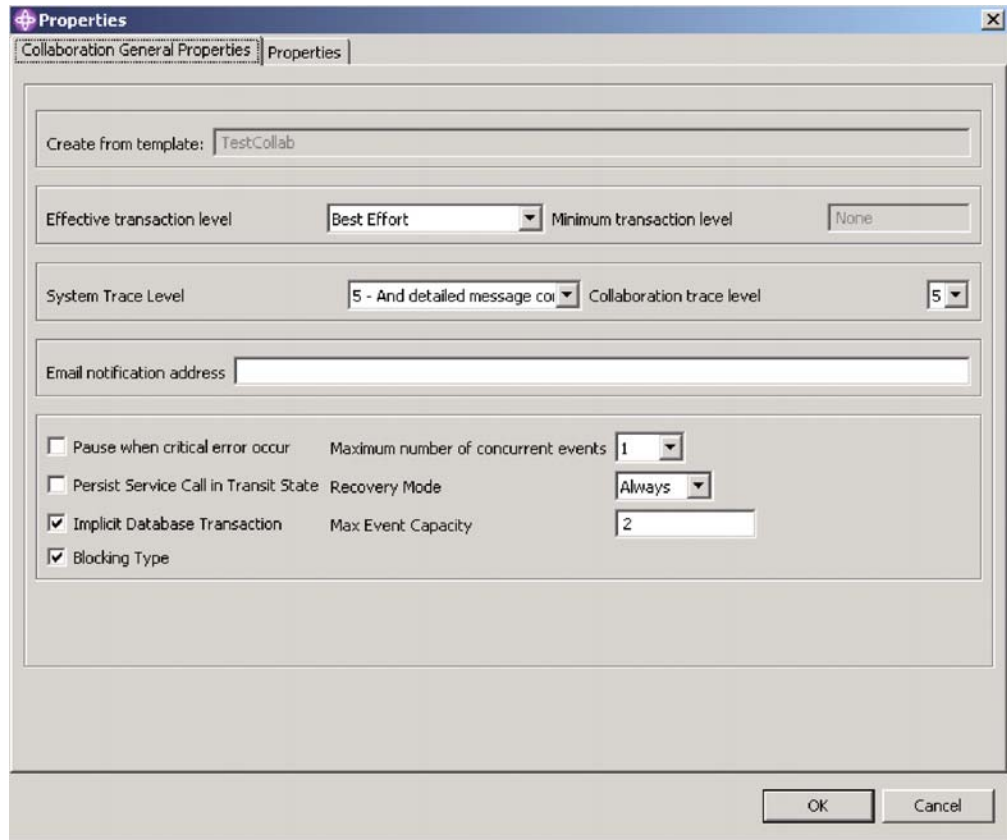


Figure 63. Properties dialog box, Collaboration General Properties tab

- In the Recovery Mode drop-down menu, choose one of the following:
  - Always
 

The collaboration will recover all WIP events whose state is WORKING and which it owns at the time of server boot.
  - Deferred
 

The collaboration will change the WIP event states to a deferred recovery state. You will then need to handle those events at a later time using the Flow Manager. For instructions on using Flow Manager, see “Using Flow Manager” on page 152.

### Asynchronous recovery

InterChange Server does not wait for collaborations and connectors to recover before it completes boot-up; collaborations and connectors are allowed to recover asynchronously after InterChange Server has booted. This makes it possible to use troubleshooting tools, such as System Monitor and Flow Manager, while the connectors and collaborations are still recovering.

## Critical errors

Critical errors in the WebSphere ICS system can cause problems in your runtime environment. A critical error as defined in the WebSphere ICS system can be generated by one of the following situations:

- Application time-out

- Inability to log on to the application
- Connector agent whose status is unknown

By default, a collaboration continues processing subsequent initiators after a flow has failed. However, a collaboration's behavior can be configured to pause automatically when a critical error occurs that could cause flows to fail. Configuring a collaboration in this way eliminates the possibility of the next flow failing for the same reason by not processing any more initiators after a flow fails. This is critical if the sequence in which initiators are processed needs to be maintained. If the collaboration pauses, the order in which initiators arrived to the server is maintained. At this point, you can fix the critical error, resolve the failed flow, then restart the collaboration. If collaborations are not dependent on an initiator that is associated with a failed flow, you can choose to resume the collaboration and resolve the failed flow at a later time. See "Failed flows" on page 151 for more information on submitting failed events.

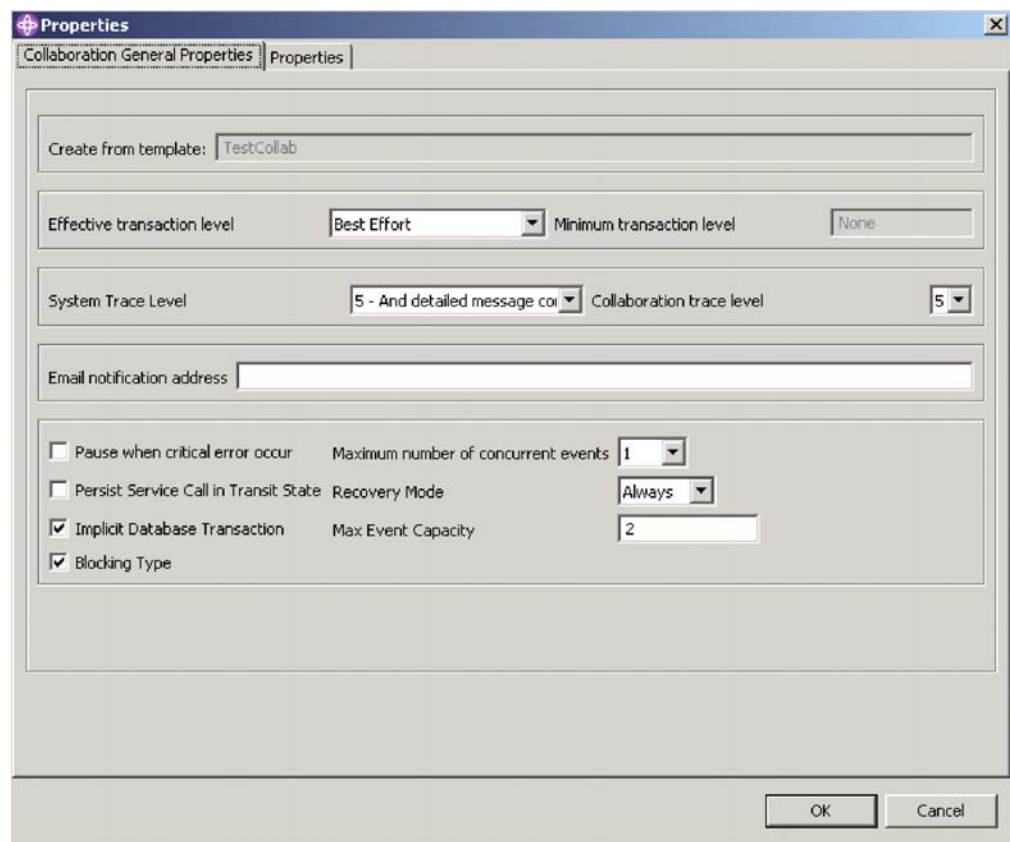


Figure 64. Properties dialog box showing "Pause when critical error occurs" option

To configure a collaboration object to pause after a critical error occurs, put a check in the "Pause when critical error occurs" box in the Collaboration General Properties tab of the Properties dialog box.

If this value is set, the collaboration will pause when a critical error occurs and will remain paused until either of the following occurs:

- The connector agent boots up again, automatically notifying the collaboration to resume processing, or
- You manually restart the collaboration, using the Unresolved Flows window

If you do not configure the collaboration to pause when a critical error occurs, the following situation might happen:

Two initiators, E1 and E2, are waiting to be processed by a collaboration. E1 creates a new customer and E2 updates E1. Since E2 updates E1, E1 must process before E2. If a critical error occurs while a collaboration is processing E1 and E1 fails as a result, then E1 is moved to the resubmission queue. If the Pause when critical error occurs box is not checked, the collaboration attempts to process E2. E2 fails because it relies on the successful processing of E1.

If the collaboration property `CONVERT_UPDATE` is set to true, then E2, which updates E1, becomes a create and creates the new customer with the updated data. Data in E1 is now old and should not be manually submitted because it will overwrite data delivered by E2.

## Lost connection to application

Collaborations that are running assume that connectors have live connections to their applications. If a connector's application becomes unavailable, the connector is unable to poll the application for events and to satisfy collaboration requests.

When an application is unavailable, a connector that polls that application for events generates an error at each polling attempt. If the connector determines that the connection with the application has been lost, the connector agent terminates and returns a status to the connector controller requesting that the connector controller also terminate.

If a collaboration sends a request to a connector while the connector is up but its application is down, the request returns with a failure status to the collaboration. This happens only if the connector property `ControlStoreAndForwardMode` is set to false. The collaboration fails, logging one of the following messages: 17050, 17058, 17059, or 17060. If you receive such messages, check the status of the application.

## Unknown connector agent status

The status of the connector agent is crucial to the WebSphere ICS system because it is a starting point for application events that processes. A connector controller maintains the status of its connector agent and relays this information to System Manager.

The connector controller maintains the status of its connector agent by sending response requests to the connector agent at 15-second intervals.

If the connector agent does not respond after three consecutive checks, its status is assumed to be unknown. An unknown connector agent status might mean that the connector agent is down or if the connector agent is installed across the network, the network connection might be down.

Setting the `ControllerStoreAndForwardMode` property for the connector to true makes the connector controller wait for the connector agent to come up before delivering any pending events. Setting this property to false makes the connector controller fail collaboration requests. The failed collaboration requests are moved to the resubmission queue and can be resubmitted using Flow Manager. See "Failed flows" on page 151 for more information.



**Attention:** Collaborations relying on an agent whose status is unknown can hang if the `ControllerStoreAndForwardMode` property is set to true. Collaborations hang until the connector agent is restarted and replies to the collaboration requests.

When the `Pause when critical error occurs` box is checked for the associated collaborations, the collaborations bound to this connector pause upon receiving the unknown status of the connector agent. An error message is logged and e-mail is sent if the e-mail connector is configured.

## Database connection failures

When InterChange Server needs a database connection for one of its services but finds that the maximum number of connections are already in use, the server tries to free a connection that is idle. If the server is unsuccessful, the connection attempt fails and InterChange Server logs error 5010: Unable to find an available connection in the cache. The maximum number of connections *max-connections-value* has been reached.

If you set a constraint on the number of InterChange connections by setting the `MAX_CONNECTIONS` parameter, you should monitor error 5010 messages because a connection failure can have undesirable consequences. For example, when InterChange Server cannot obtain a connection for its event management service, it stops running. By default, this constraint is set to an unlimited number of connections.

Connection failures indicate that the maximum number of allocated connections is insufficient to meet the runtime work load. If you cannot allocate more connections to InterChange Server in the current database, consider partitioning its work load across multiple databases.

## Flow failures

At times, the WebSphere ICS system or its associated applications may fail. Successfully processing flows that carry data through the WebSphere ICS system is critical, so in a runtime environment it is critical to maintain data consistency. System failures such as system errors, data errors, and critical errors can cause flows to fail to process. The WebSphere ICS system has some built-in capabilities that allow you to handle system failures. The following topics describe two different types of failures:

“Failed flows” on page 151

“Failed transactional collaborations” on page 152

For information on managing, resolving, and preventing flow failures, see “Using Flow Manager” on page 152.

### Failed flows

A system configuration, object definition, application-specific, or data consistency error can cause a flow to fail when the WebSphere ICS system is processing that flow. Improperly functioning InterChange Server components, such as business object mapping failures, or the unavailability of a connector, can generate system errors, which cause flows to fail. Data inconsistencies, such as an isolation violation of application data during execution of a collaboration, generate data errors, which also cause flows to fail.

If an error occurs while a connector controller or a collaboration is processing a flow, the flow fails and is moved to the event resubmission queue. From here, you have the following choices:

- Submit the event with the original business object
- Submit the event using the latest business object
- Discard the event from the system

For instructions on resolving failed flows, see “Using Flow Manager” on page 152..

### **Failed transactional collaborations**

System and data errors can cause a transactional collaboration to fail. When one of these errors occurs, the collaboration attempts a rollback. If the rollback of a collaboration’s compensation steps fails, the collaboration is in an “in-doubt” state. If an error occurs during runtime recovery, the collaboration is put into a list of failed transactional collaborations owned by the corresponding collaboration. A failed transactional collaboration is a collaboration whose compensation steps failed to roll back.

Once a transactional collaboration fails, you need to resolve it. You can handle a failed transactional collaboration by using Flow Manager. For instructions on resolving failed transactional collaborations, see “Using Flow Manager” on page 152..

**Preventing failed transactional collaborations from pausing:** The default behavior for a failed transactional collaboration is to pause. You can prevent failed transactional collaborations from pausing by adding a property called `PAUSE_ON_COMPENSATION_FAILURE` to the collaboration template and changing the setting from `TRUE` (default) to `FALSE`. To add the new property and change the setting to `FALSE`, do the following:

1. In System Manager, double-click the collaboration template that failed. Process Designer opens.
2. Double-click the Definitions icon in the left-hand frame. The Template Definitions window appears in the right-hand frame.
3. In the Properties tab, click Add. A Name dialog box appears.
4. Type `PAUSE_ON_COMPENSATION_FAILURE` in the Name field, then click OK. The `PAUSE_ON_COMPENSATION_FAILURE` property appears under the General Properties node in the left-hand pane.
5. With `PAUSE_ON_COMPENSATION_FAILURE` selected in the left-hand pane, select Boolean from the Property Type drop-down menu.
6. Uncheck the `IsDefaultVal` box for the “true” row, and place a check in the `IsDefaultVal` box in the “false” row.
7. Click Apply.
8. Close Process Designer.

---

## **Using Flow Manager**

To locate, view, and handle failed events, use Flow Manager, a tool that is installed with the IBM WebSphere ICS product. Flow Manager allows you to easily construct a query to locate and display unresolved flows. After you display the unresolved flows, you can select any flow in the display and choose to submit it, discard it, or perform other actions. For background information on failed flows, see “Flow failures” on page 151.

The following topics describe how to use Flow Manager for constructing the queries, viewing the details, and handling the events:

“Starting Flow Manager”

“Finding unresolved flows” on page 155

“Viewing details for unresolved flows” on page 158

“Handling unresolved flows” on page 160

## Starting Flow Manager

To start Flow Manager, do the following:

1. Click Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration > Administrative > Flow Manager. The Connect to InterChange Server dialog box appears (see Figure 65).



Figure 65. InterChange Server dialog

2. Enter the server name, using any one of the following methods:
  - Type the name of the server in the Server Name field.
  - Select the server name from the Server Name list.
  - Click the Browse button to browse for the server on the network. The following figure shows the Server dialog that opens when you click this button.

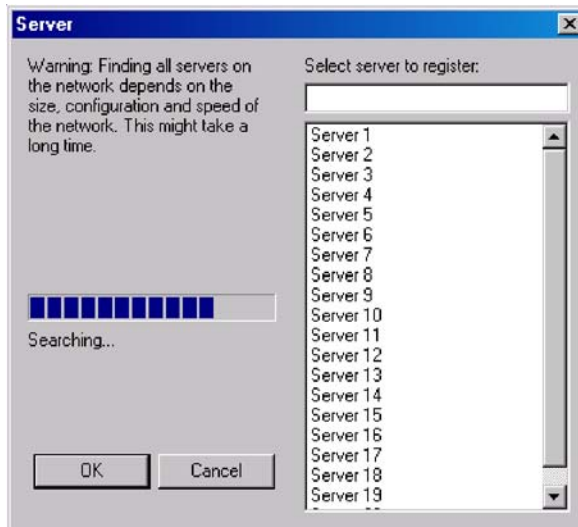


Figure 66. Server dialog

**Note:** Browsing for the server on the network could take a long time, depending on how many servers exist on the network.

3. Enter the user name and password for the server you want to connect to. If you check the "Remember user name and password" checkbox, the user name and password will be stored in the registry along with the server name, and the password will be encrypted. Click Connect.

**Note:** To remove unwanted users or servers from the registry, click Options. This opens the Options dialog, from which you can remove servers or users.

The following figure shows the Options dialog.

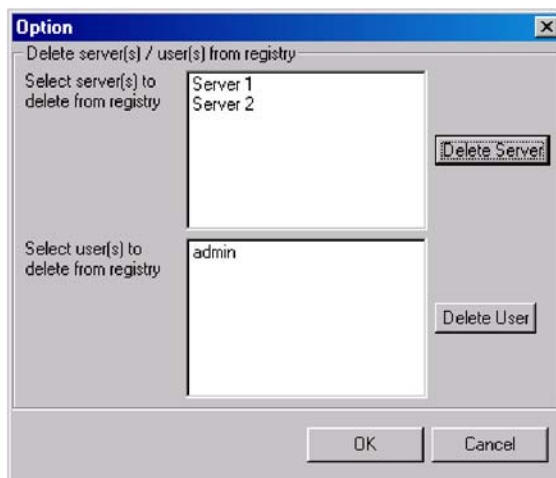


Figure 67. Options dialog

After clicking Connect, Flow Manager opens (see Figure 68 on page 155).

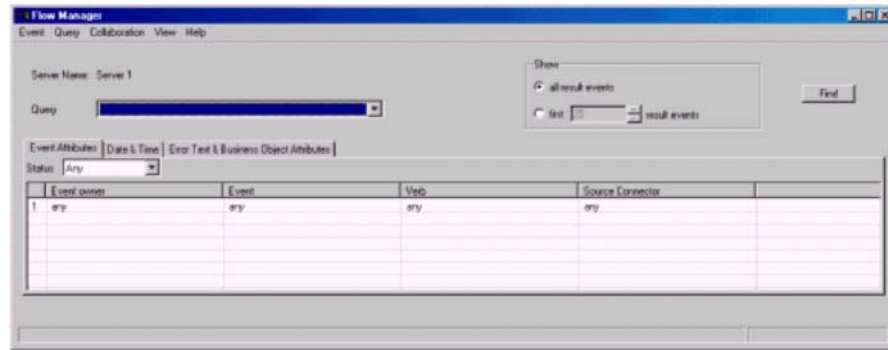


Figure 68. Flow Manager

## Finding unresolved flows

You can construct queries that select only the unresolved flows that match criteria that you specify. To construct queries, you must first start Flow Manager. See “Starting Flow Manager” on page 153 for instructions on starting Flow Manager. The following topics describe how to construct queries to find unresolved flows.

“Finding all unresolved flows”

“Finding specific unresolved flows”

“Managing queries” on page 158

### Finding all unresolved flows

If you want to locate and display all unresolved flows, accept the defaults of “any” for all the fields, then click the Find button or select Find from the Event menu. The system will locate all unresolved flows and display their details.

### Finding specific unresolved flows

If you want to construct a query to find specific flows, refer to the following sections that describe the search options in each tab. After specifying your search criteria, click the Find button or select Find from the Event menu.

- “Event Attributes tab” on page 155
- “Date & Time tab” on page 156
- “Error Text & Business Object Attributes tab” on page 157

**Event Attributes tab:** Choose the Event Attributes tab to search according to the characteristics of the unresolved event (see Figure 69 on page 156).




Figure 69. Event Attribute tab

The following options are available from the Event Attributes tab:

- Status: select events according to their status:
  - Any: Selects all unresolved events
  - Failed: Selects all unresolved events that are in the Failed state.
  - In Transit: Selects all unresolved events that are in the In Transit state.
  - Possible Duplicate: Selects all unresolved events that are in the Possible Duplicate state
  - Deferred: Selects all unresolved events in which the point of failure was a collaboration with a recovery mode setting of Deferred.
  - Waiting: Selects all unresolved events in which the point of failure was a collaboration with a recovery mode setting of Waiting.
- Event Owner
 

Select a collaboration name or a connector name to query only for flows that failed within that collaboration or connector. You can string together multiple selections by clicking in the new empty field that is created below each existing selection. To search all of you collaborations and connectors, choose Any.
- Event
 

For each collaboration or connector that you select, choose an event, as represented by a business object. You can choose one event, string together multiple events, or choose Any.
- Verb
 

Chose the verb to query for each event.
- Source Connector
 

Choose the source connector or connector whose flows you want to query, or choose Any.

**Date & Time tab:** Use the Date & Time tab to query only the flows that failed during a specific time period (see Figure 70 on page 157).

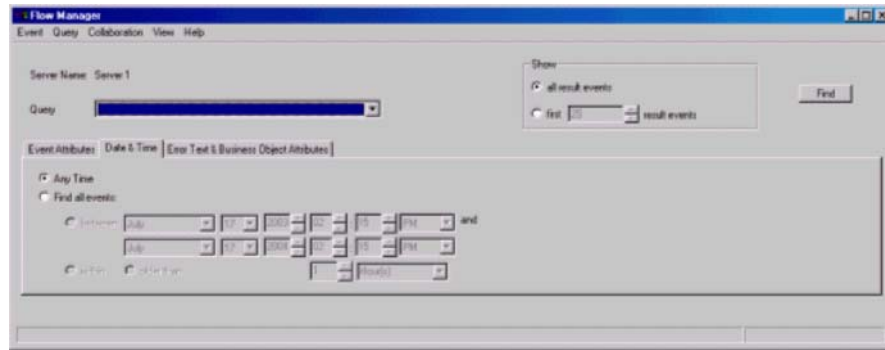


Figure 70. Date & Time tab

The following options are available from the Date & Time tab:

- Any Time: Queries for all failed events, with no restriction on when they may have occurred
- Find All Events: Queries for all failed events that occurred within specified time periods:
  - between (specified times)
  - within (period of days)
  - older than (period of days)

**Error Text & Business Object Attributes tab:** Use the Error Text & Business Object Attributes tab to construct a query that includes only certain business objects, business object attributes, and attribute values (see Figure 71).

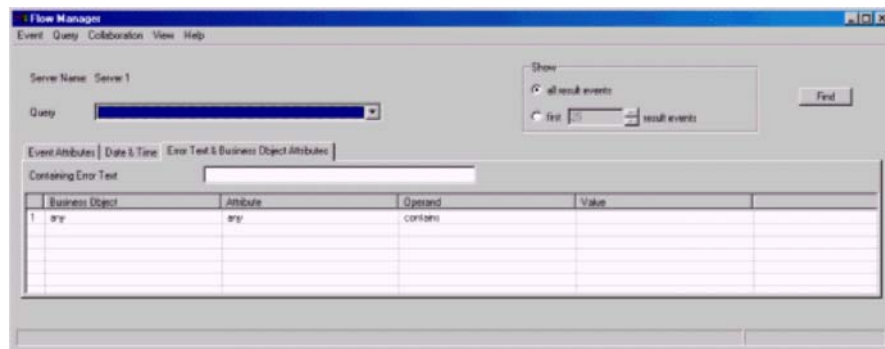


Figure 71. Error Text & Business Object Attributes tab

Make selections in the following columns:

- Business Object
- Attribute
- Operand
- Value

You can also query only unresolved flows that produce a message containing specified text that you enter in the Containing Error Text field.

**Note:** The key words you enter in the Containing Error Text field are not case-sensitive.

## Managing queries

You can use the menu options in the following table to manage queries. For information on taking action on unresolved flows, refer to “Handling unresolved flows” on page 160.

Table 27. Flow Manager menu options

Menu option	Description
Event > Find	Finds all the unresolved events with the specified filter information. As a shortcut, you can also click the Find button.
Event > Delete	Deletes one or more selected events in the result frame.
Event > Save	Saves all the events in the result frame in an Excel format.
Event > Print	Prints the selected events or all the events in the result frame in a table format. A dialog appears prompting you to select Print All Events or Print Selected Events.
Event > Get Count	Gets a number of the events corresponding to the current query.
Event > New Search	Clears all the result events, close the result frame, and refresh all the filter controls.
Event >Exit	Exits Flow Manager.
Query > Delete	Deletes the query that appears in the Query menu.
Query > Save	Save the current filter information as a query; the query name entered appears in the Query menu.
Query > Show/Hide	Show or hide the content of the query.

## Viewing details for unresolved flows

To view details for unresolved flows, select a query from the Query menu, then click the Find button or select Find from the Event menu. The results appear in the bottom half of the Flow Manager window (see Figure 72 on page 159).



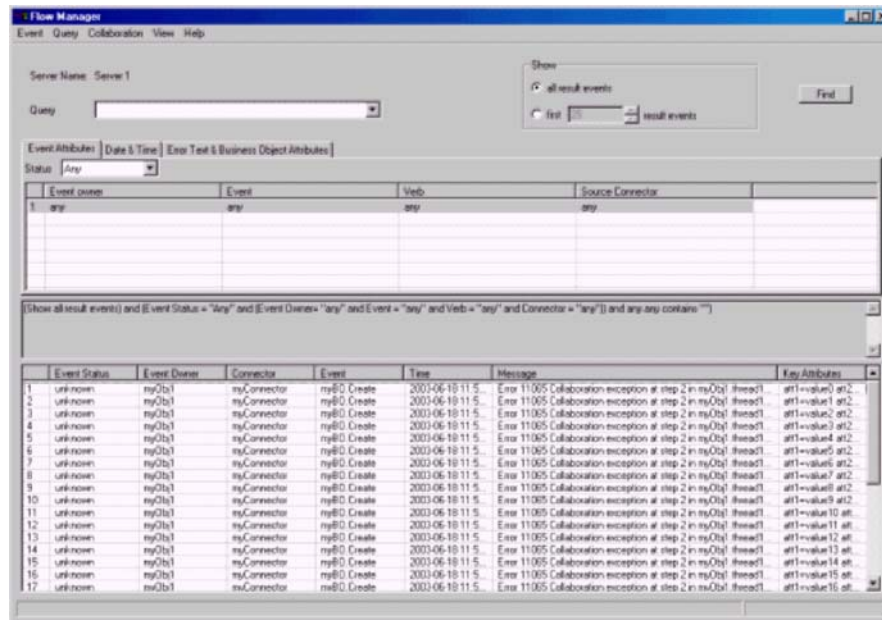


Figure 72. Flow Manager displaying query results

The system displays an unresolved flows table, containing a list of events with the following information:

- Event Status
- Event Owner
- Point of Failure
- Connector
- Event
- Time
- Message
- Key Attributes

For the waiting events, the Event Status cell contains the following information:

Waiting  
[timeout expiration:...]  
[scenario name:...]  
[scenario node ID...]

To access more information about any of the events, double-click the row containing the event. This opens the Show Event Details dialog (see Figure 73 on page 160).

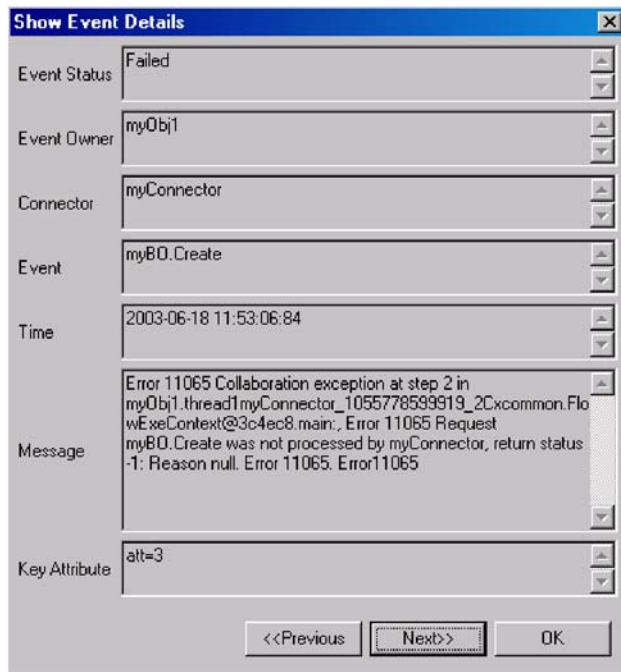


Figure 73. Show Event Details dialog

If you want to take action on any of the unresolved flows, refer to the next section, "Handling unresolved flows."

## Handling unresolved flows

You can take an action on any unresolved flows in the results section of Flow Manager by using either the Events menu or the menu that appears when you right-click the event.

The following table describes both the Events menu options and the right-click menu options available for each event listed in the results section of Flow Manager.

Table 28. Actions you can perform on unresolved events

Event menu option	Right-click menu option	Description
Event > Delete	Right-click > Delete	Deletes one or more selected events in the result frame.
Event > Save	Right-click > Save	Saves all the events in the result frame in an Excel format.
Event > Submit	Right-click > Submit	Sends the data of the events back to the destination application.
Event > Refresh & Submit	Right-click > Refresh Submit	Refreshes then submits the selected events.
Events > Cancel Waiting	Right-click > Cancel Waiting	Cancels the waiting events (long-lived business processes) in the results section.

Table 28. Actions you can perform on unresolved events (continued)

Event menu option	Right-click menu option	Description
Event > Display Details	Right-click > Display Details	Displays a "Show business object data" dialog which contains detailed information about the business object associated with the selected event, including business object definition, verb, and attributes. You can browse the display details of the entire list of events in the results section by clicking Next or Previous.
Event > Select All	Right-click > Show Event Details	Displays the Show Event Details dialog. This dialog contains details about the selected event, including the full text in the Message field. This is useful because the message is only partially viewable from the results section of Flow Manager. You can browse the event details of the entire list of events in the results section by clicking Next or Previous.
Event > Print	Right-click > Print	Prints the selected events or all the events in the result frame in a table format. A dialog appears prompting you to select Print All Events or Print Selected Events.

Figure 73 on page 160 shows the Show Event Details dialog that appears when you right-click an event then select Show Event Details.

Figure 74 on page 162 shows the Business Object Data dialog that appears when you right-click an event then select Display Details.

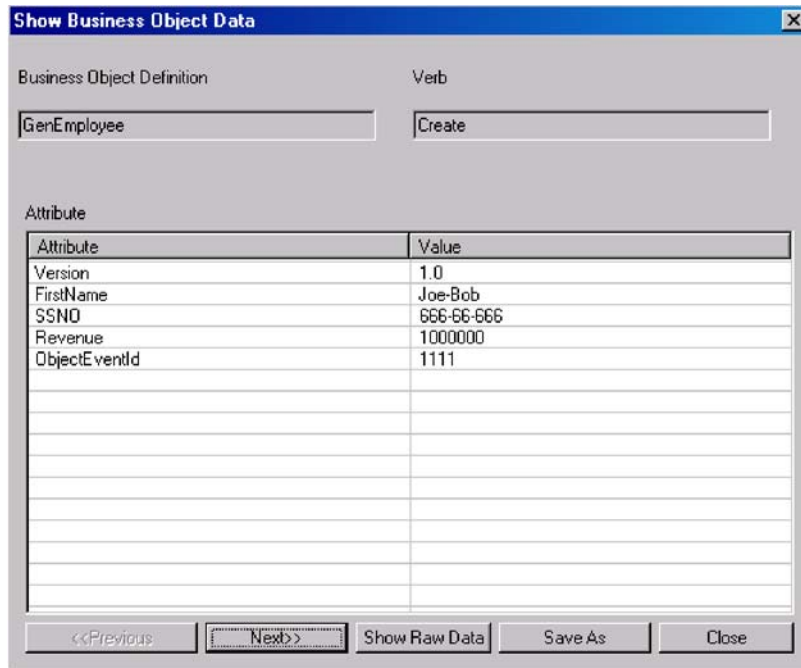


Figure 74. Show Business Object Data dialog

## Resolving collaboration deadlocks

A deadlock is a situation where two or more processes are unable to proceed because each is waiting for the other processes to proceed. Deadlocks are an undesirable side effect of the concurrency control provided by event isolation within a collaboration. See the *Collaboration Development Guide* for more information on event isolation.

Figure 75 on page 163 illustrates a deadlock between two active collaboration groups resulting from the following sequence of events:

1. At time T1, Collaboration A1 receives an event, E1, then makes a service call to Collaboration B2 and sends a child business object for E1. Collaboration A1 waits for the service call to complete.
2. At time T2, Collaboration B1 receives an event, E2, then makes a service call to Collaboration A2 and sends a child business object for E2. Collaboration A2 waits for the service call to complete.
3. At time T3, Collaboration B2 is waiting for Collaboration B1, since B2 and B1 have the same port binding and the event from B1 was delivered before the event for B2 arrived.
4. At time T4, Collaboration A2 is waiting for Collaboration A1, since A2 and A1 have the same port bindings and the event from A1 was delivered before the event for B1 arrived.

At this point, all collaborations are unable to move forward.

**Note:** A port binding consists of the business object type and connector name. See the *Collaboration Development Guide* for more information on port bindings.

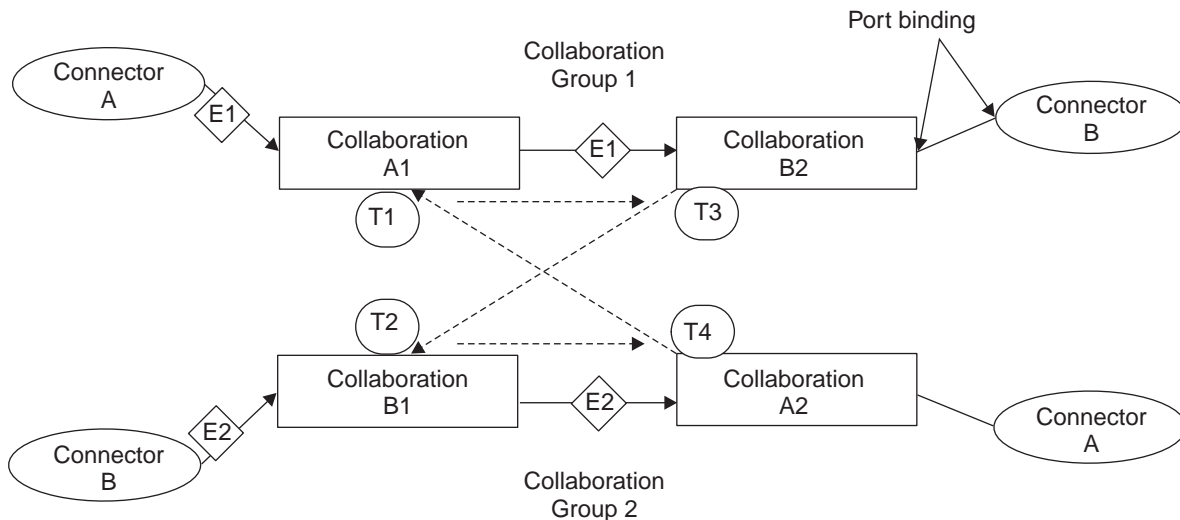


Figure 75. Deadlock between collaboration groups

This section covers the following topics:

“Detecting a collaboration deadlock” on page 163

“Detecting group collaboration deadlocks” on page 164

“Fixing a collaboration deadlock” on page 164

## Detecting a collaboration deadlock

You can configure the IBM WebSphere ICS system to either perform or not perform deadlock detection:

- By default, the IBM WebSphere ICS system performs deadlock detection automatically when you start InterChange Server. However, when deadlock detection is performed, startup of InterChange Server can be delayed if a collaboration group contains many collaboration objects, because ICS must traverse all the collaboration objects in the group to determine whether there is a deadlock in the group. This can cause slow startup even when no deadlock exists
- You can configure the IBM WebSphere ICS system to not perform deadlock detection. If you do so, the IBM WebSphere ICS system will start collaboration groups without first checking for deadlocks. This can make it possible for InterChange Server to boot more quickly. However, if deadlock detection is not performed and a deadlock exists, an event that is later sent to a collaboration may fail.

System Manager does not provide the ability to set the DEADLOCK\_DETECTOR\_CHECK configuration parameter. Instead, to set this configuration parameter, you must edit the InterchangeSystem.cfg file and change the parameter’s value in this file. Because the InterchangeSystem.cfg file is an .xml file, the following lines should exist in this file to define the DEADLOCK\_DETECTOR\_CHECK parameter:

```
<tns:name>DEADLOCK_DETECTOR_CHECK</tns:name>
<tns:value xml:space="preserve">>false</tns:name>
```

To restore deadlock detection, change the false value to true.

## Detecting group collaboration deadlocks

You can check for a group collaboration deadlock in one of the following ways:

- In System Manager, right-click the running group collaboration, then select Diagnostics.

A window appears with the following message:

The following diagnostic tests were run on this collaboration:

This message is followed by one of the following results:

- Deadlock detection - Fail
- Deadlock detection - Pass
- No diagnostic tests supported at this time

**Note:** This result appears only if you run Diagnostics on a non-group collaboration.

- Check the InterchangeSystem.log file for the following error at the time the hanging collaborations were started:

Error 11135: Activation of collaboration *collaboration\_name* group could cause a potential deadlock with one or more existing collaboration groups, and is therefore disallowed.

This error warns only of a potential deadlock situation. The informational messages preceding error 11135 identify the active collaboration groups that potentially enter into a deadlock.

## Fixing a collaboration deadlock

If the WebSphere ICS system encounters a deadlock, you must shut down and restart InterChange Server. First, gracefully shut down all other collaborations, then shut down the server immediately. See “Shutting down InterChange Server” on page 49 for more information on shutting down the WebSphere ICS system.

Upon system restart, a hung collaboration that caused the deadlock automatically starts and resubscribes to all of the business objects it supports. The business objects that caused the collaborations to enter into a deadlock are redelivered. At this time, the collaborations should not enter into another deadlock because deadlocks are timing dependent. It is unlikely that you will have the exact same server load and isolation sequencing that you had when your system encountered the deadlock.

After restarting the system, shut down the collaborations involved and rebind the ports so that this does not occur again.

## Preventing collaboration deadlocks

You can prevent collaboration deadlocks by configuring the deadlock retry settings in the Database tab of server configuration screen in System Manager. To configure the deadlock retry mechanism, do the following:

1. From System Manager, right-click the server under Server Instances, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool in which you can edit the InterchangeSystem.cfg file.
2. Click the Database tab. A dialog box appears in the upper-right section of the System Manager window in which you can enter the parameters necessary for configuring database configuration at the system level (see Figure 76 on page 165).

3. In the "Max database retry" field, enter a number that represents the maximum number of retries you want the server to perform if a deadlock occurs.
4. In the "Deadlock retry interval" field, enter a number that represents the number of seconds you want the system to wait before retrying.

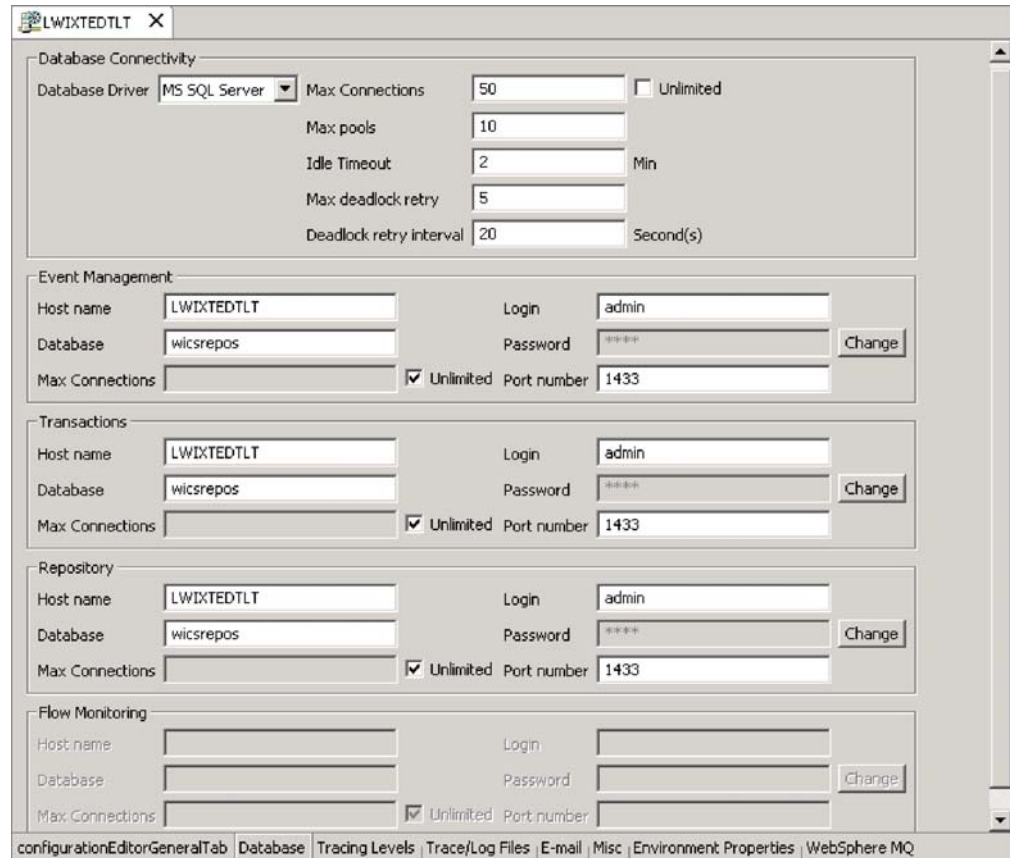


Figure 76. Edit Configuration screen, Database tab

## Managing WebSphere MQ

This section contains information about resolving problems related to WebSphere MQ. The following topics are covered:

“Errors related to queue configuration” on page 165

“Failures to create new instances of MQQueueManager” on page 166

“Failure recovery and the queue manager” on page 166

“Memory checker segmentation fault” on page 166

### Errors related to queue configuration

WebSphere MQ may report errors if the configuration settings for the queues do not allow for a sufficient number of messages, a sufficient maximum length for messages, or a sufficient size for log files. For information about adjusting these configuration settings, see “Configuring WebSphere MQ Message Queues” in the *System Installation Guide*.

## Failures to create new instances of MQQueueManager

On the AIX platform, the use of the bindings transport (C via JNI) in JDK 1.3.1\_06 requires some additional configuration to avoid a conflict with shared memory segments. If the configuration is not done, a failure to create a new instance of MQQueueManager, with reason code 2059, can occur. The failure can be accompanied by the generation of FDCs in the /var/mqm/errors directory.

To perform the configuration, use an additional parameter for the mqs.ini file. The parameter is IPCCBaseAddress and is set on a per queue manager basis. By default this parameter is set to the value 8, but it is recommended that you set this value to 12. The following is an example of an altered QueueManager stanza in an mqs.ini file:

```
QueueManager:  
  Name=MQJavaTest  
  Prefix=/var/mqm  
  IPCCBaseAddress=12
```

## Failure recovery and the queue manager

After a full system failure and recovery, the queue manager might fail to start, displaying this message:

```
AMQ7017 Log not available.
```

This typically indicates that the log file is missing or damaged, or that the log path to the queue manager is inaccessible.

To remedy the problem, create a dummy queue manager and copy its header log file and transaction log files over to the actual queue manager. This will allow you to restart the queue manager.

## Memory checker segmentation fault

The memory checker thread can get a segmentation fault when JMS or other modules are used in WebSphere MQ binding mode.

By default, Java 1.3.1 runs using eight segments. (Eight is the maximum value allowed; each segment is 256 MB). You can reduce the number of segments Java uses with the LDR\_CNTRL environment variable. For example, if you want Java to use only three segments, set LDR\_CNTRL to set the following value prior to starting Java:

```
export LDR_CNTRL=MAXDATA=0x30000000
```

If your native methods use a large amount of shared memory (for example, with functions like `shmat()` and `mmap()`), you might need to reduce the number of segments used by Java. You might also find that you need to reduce the number of segments used by Java when concurrently running Java with other applications that use shared memory, particularly those that explicitly specify the address where a shared memory segment is attached (rather than letting AIX choose the address).

---

## Managing database connection pools

At run-time, it takes time for processes to establish new connections to the database. You can minimize that time by establishing database connection pools for use by your collaboration and map processes.



Use of database connection pools requires establishing pool names in collaboration template and map code, in addition to creating the database connection pools in System Manager. For instructions on performing those tasks, and additional information about setting up database connection pools, see the *System Implementation Guide*.

If database connection pools have already been set up on your system, you can perform the following tasks to manage how the database connection pool resources are allocated to your collaboration and map processes:

- Revise the number of connection processes that can be used by a specific database connection manager and its pools
- Revise the database connection pool assignments.

For details about these tasks, see the *System Implementation Guide*.

---

## Managing High-Availability (HA) systems

A high-availability (HA) system is made up of two or more machines (the primary and one or more backup machines) that are configured identically and designated as a cluster. Each machine is considered a node in the cluster. The primary and backup nodes share a cluster name and IP address. External processes use this name and IP address to access a service on the cluster, which runs on either the primary or one of the backup nodes. All nodes have access to a shared Redundant Array of Independent Disks (RAID) storage system. For windows systems only, the shared RAID storage system is used only by the active node.

The HA configuration provides shutdown and automatic restart of unresponsive (failed) software programs, and migration to the cluster backup node when failures on the active node are detected. The cluster backup node assumes the cluster name and IP address and automatically takes over system processing until such time as the failure is corrected on the primary node and a failback is initiated (that is, manually return processing to the original system).

This section provides the following information about how to manage a high-availability system that includes InterChange Server:

“Supported HA environments”

“Maintaining a Windows HA system” on page 168

### Supported HA environments

The high-availability (HA) option is available on the following operating systems:

- UNIX-based systems -- Sample scripts with a README file are provided as part of an IBM SupportPac.
- Windows 2000 -- *The System Installation Guide for Windows* provides a set of instructions on how to set up HA using Microsoft Cluster Server (MSCS) software. An IBM Category 2 SupportPac provides some dynamic link library (.dll) files for use with MSCS.

**Note:** Scripts and files that assist with the HA implementation are available as unsupported in Category 2 SupportPac. You can find the information about these IBM SupportPacs at:

<http://www.ibm.com/software/integration/supportpace/category.html#cat2>

In addition, both the *System Installation Guide for UNIX* and *for Windows* provide basic instructions on how to configure the hardware and software for use in an HA environment.

## Maintaining a Windows HA system

Once an HA system is set up according to the configuration instructions provided in the *System Installation Guide for Windows*, it should need minimal maintenance or reconfiguration. This section summarizes some of the tasks for maintaining an HA system that is set up on a Windows operating system to use the Microsoft Cluster Server (MSCS) software. The following topics are covered:

- “Checking cluster status”
- “Detecting a failover”
- “Moving groups to perform maintenance”
- “Changing the status of a resource” on page 169

### Checking cluster status

The MSCS Cluster Administrator is the primary administration tool that you use to administer and check the status of the cluster. Each resource, such as InterChange Server, is listed along with its state (online or offline, failed, or online or offline pending), owner (node 1 or node 2), and the type of resource (a description such as InterChange Server, disk resource, or connector). From this window, you can see the status of the individual services and which of the cluster nodes is active.

Other Windows administrative tools provide information about the status of the cluster. In particular, check the MSCS online help and documentation for details about using the following tools to monitor the cluster:

Windows Event Viewer	View and manage System, Security, and Application event logs
Windows Services option in the Control panel	Verify that the Cluster Service is running

### Detecting a failover

Several types of icons appear in MSCS Administrator beside the various listings of groups, resources, and other elements. The most important icon to recognize is the node-down icon, which indicates that a failure has occurred on a cluster node and that its groups and resources have been transferred to the surviving cluster node. A node-down icon displays in the MSCS Administrator as a cluster node icon with a red x through it.

The node-down icon does not necessarily mean that you have lost functionality in any of your groups or resources. In normal operation, the group fails over to the backup cluster node.

### Moving groups to perform maintenance

When you stop the Cluster Service on a node, you prevent clients from accessing cluster resources through that node and all groups move to the other node (if the failover policies allow it). This can be useful when you need to take the primary node offline to perform maintenance or upgrade its software. The following steps describe how to perform a move to the backup node for maintaining or upgrading the primary node:

1. Stop the Cluster Service on the backup node by clicking its node icon, clicking the File menu, and selecting Stop Cluster Service.
2. Upgrade the backup node if a software upgrade is the maintenance being performed. Be sure to place executables and libraries on the node disk and data files on the shared RAID.
3. Restart the Cluster Service on the backup node by clicking File > Start Cluster Service.
4. On the primary node, right-click each group and select Move Group. This allows you to take a server offline without losing availability of your resources.
5. Stop the Cluster Service on the primary node by clicking its node icon, then clicking File > Stop Cluster Service.
6. Upgrade the primary node in the same manner as you upgraded the backup node.
7. Restart the Cluster Service on the primary node by clicking File > Start Cluster Service.
8. On the backup node, right-click each group and select Move Group to move the groups back to the primary node.

### Changing the status of a resource

Use Cluster Administrator to manually bring individual resources online or take them offline. Change the status by selecting the resource, and from the File menu, selecting either Bring Online, Take Offline, or Initiate Failure. You should only off or online WebSphere MQ, connectors, and InterChange Server.

---

## Managing runtime properties

The components of the WebSphere InterChange Server system obtain much of their initialization information from a single global environment file:

### UNIX

```
ProductDir/bin/CWSharedEnv.sh
```

### Windows

```
ProductDir\bin\CWSharedEnv.bat
```

The startup script of the various components in the InterChange Server system read this file as part of their initialization process:

- An adapter
- InterChange Server
- SNMP agent
- repos\_copy utility
- Many InterChange Server tools

This CWSharedEnv file contains initialization information for the communication and runtime software that the InterChange Server system uses. Most of this information is provided in the CWSharedEnv file that comes with the product and never needs to change. Some information is customized for your system during installation.

If you need to change runtime information *after* installation, you can take either of the following approaches:

- Edit the appropriate variable in the CWSHaredEnv file. When you change the property in the CWSHaredInv file, you change it for *all* components that read this file during their setup process.
- Specify the property and its value as a command-line option to the component's startup script. When you change the property on the command line, you change it only for the component that you are starting. Properties you specify on the command line override any other property settings within the system or from the CWSHaredEnv file.

**Note:** For more information on how to change the runtime properties that the IBM Java ORB supports, see "Customizing the ORB."

---

## Managing the Object Request Broker

This section contains information about resolving problems related to the IBM Java Object Request Broker (ORB), which handles communication between InterChange Server and several of its components. This section provides information on the following topics:

- "Installing the ORB"
- "Customizing the ORB"
- "Changing the location of the ORB" on page 172
- "Using the IBM Transient Name Server" on page 173

### Installing the ORB

InterChange Server requires the IBM Java ORB to communicate with several of its components, including adapters, System Manager, and SNMP agents. Use of the ORB requires installation of the IBM Java ORB. The IBM Java ORB is installed as part of the IBM Java Runtime Environment (JRE) software, which the WebSphere InterChange Server Installer installs automatically. These steps and how to configure the ORB for various environments are described in the *System Installation Guide for Windows*.

### Customizing the ORB

Table 29 lists the properties that the IBM Java ORB supports to customize its behavior.

*Table 29. IBM Java ORB properties that can be customized*

IBM ORB property	Configuration parameter	Description
com.ibm.CORBA.ListenerPort	0Aport	Port number on which the ORB server (within InterChange Server) listens for incoming requests.
com.ibm.CORBA.LocalHost	0AipAddr	IP address or host name of the machine on which the ORB server (within InterChange Server) is running.

Table 29. IBM Java ORB properties that can be customized (continued)

IBM ORB property	Configuration parameter	Description
com.ibm.CORBA.ThreadPool.MaximumSize	OAThreadMax	Maximum number of threads that the connection manager can create. The default value (zero) indicates that no size restriction exists.
com.ibm.CORBA.ThreadPool.InactivityTimeout	OAThreadMaxIdle	The time (in seconds) before an idle thread will be destroyed.
com.ibm.CORBA.RequestTimeout	None	Number of seconds that a CORBA request waits before timing out. By default, there is no timeout; the ORB waits indefinitely for a response.
com.ibm.CORBA.LocateRequest	None	Timeout value (in seconds) for Locate Requests.
com.ibm.CORBA.FragmentTimeout	None	Maximum length of time that the ORB waits for second and subsequent message fragments before it times out. Set this property to zero to indicate no timeout should occur. The default value is 30000.

You can specify these properties in any of the following ways:

- Specify the IBM ORB property and its value as a command-line option to the component's startup script. When you specify the IBM ORB property on the command line, you change it only for the component that you are starting. You specify an ORB property by preceding it with the `-D` command-line option. Properties you specify on the command line override any other property settings within the system or from the `CWSharedEnv` file.
- Edit the appropriate variable in the `CWSharedEnv` file. When you change the variable in the `CWSharedEnv` file, you change it for *all* components that read this file during their startup process. These components include any adapters, SNMP agents, InterChange Server instances, the `repos_copy` utility, and the InterChange Server tools.

**Note:** For more information on the `CWSharedEnv` file, see "Managing runtime properties" on page 169.

As Table 29 on page 170 shows, the ORB location is specified by special variables in the `CWSharedEnv` file. You must modify these variables in the `CWSharedEnv` file to change its location.

Other ORB properties are listed in the `ORB_PROPERTY` variable of the `CWSharedEnv` file. In this variable, each IBM ORB property is preceded by the `-D` command-line option. To add or change an ORB property, you must add or change the appropriate `-D` option in the `ORB_PROPERTY` variable of the `CWSharedEnv` file. Properties you specify in the `CWSharedEnv` file override any other settings from the configuration file.

- Specify the configuration parameter (if one exists) in the appropriate configuration (`.cfg`) file. You can set many of the ORB properties with configuration parameters in the CORBA section of the configuration file. Both the ICS configuration file (`InterchangeSystem.cfg`) and the adapter local

configuration file can contain a CORBA section. When you specify the configuration parameter in the CORBA configuration file, you change it for all ORB-related tasks that the ORB server performs.

**Important:** The configuration files are in XML format. Do *not* modify these files unless you use an XML editor or are very familiar with XML format!

For example, to specify the maximum number of threads, you can take any of the following actions:

- Add the IBM ORB property to the ORB\_PROPERTY variable in the CWSHaredEnv file:

```
ORB_PROPERTY=-DORBNamingProvider=CosNaming
-Dorg.omg.CORBA.ORBClass=com.ibm.CORBA.iiop.ORB
-Dorg.omg.CORBA.ORBInitialPort=%ORB_PORT%
-Dorg.omg.CORBA.ORBInitialHost=%ORB_HOST%
-Dcom.ibm.CORBA.Debug.Output=null
-Dcom.ibm.CORBA.ThreadPool.MaximumSize=100
```

- Specify the IBM ORB property on the command line when you start up the component:

```
start_server.....-Dcom.ibm.CORBA.ThreadPool.MaximumSize=100
```

- Add the OAtHreadMax configuration parameter to the CORBA section of the configuration file:

```
[CORBA]
OAtHreadMax=100
```

## Changing the location of the ORB

Table 30 lists the default ORB-location information. You can change this default information during installation. In its Naming Server screen, the WebSphere InterChange Server Installer prompts you for the IP address and port number for the InterChange Server instance. The installer saves this information in the appropriate variables of the product directory. Table 30 also shows the variables within the CWSHaredEnv file that specify the ORB location.

Table 30. Location of the IBM Java ORB

ORB location	IBM ORB property	Default value	CWSHaredEnv variable
IP address	org.omg.CORBA.ORBInitialHost	Name of the local host	ORB_HOST
Port number	org.omg.CORBA.ORBInitialPort	14500	ORB_PORT

If you need to change the ORB-location information *after* installation, you can take either of the following steps:

- Edit the appropriate variable (from the column in Table 30 labelled "CWSHaredEnv variable") in the CWSHaredEnv file. Changing the ORB location within the CWSHaredEnv file means that you change the ORB location for all invocations of all startup scripts that use the CWSHaredEnv file. For example, to change the port number of the ORB to 15002, you can set the ORB\_PORT property in the CWSHaredEnv file, as follows:  
set ORB\_PORT=15002

#### UNIX

On UNIX systems, any value you assign to the ORB\_PORT variable in the CWSharedEnv.sh file will be overridden by the user's setting of the ORB\_NAMING\_PORT environment variable (if that environment variable is set). Instead of changing ORB\_PORT within the CWSharedEnv.sh file, you can instruct all users to set the ORB\_NAMING\_PORT environment variable to the desired value.

- Specify the appropriate IBM ORB property (from the column of Table 30 on page 172 labelled "IBM ORB Property" and its value in the command line of the component's startup script with the -D command-line option. Changing the ORB location on the command line of a startup script means that you can change the ORB location only for that invocation of the component that the startup script is starting. For example, to change the port number of the ORB to 15002, you can specify the following -D option on the startup script's command line:  
-Dorg.omg.CORBA.ORBInitialProt=15002

For information on how to change ORB properties, see "Customizing the ORB" on page 170.

## Using the IBM Transient Name Server

Using IBM Transient Naming Server (tnameserv) provides the naming service for the WebSphere business integration system. When a component of the WebSphere business integration system starts, it registers itself with the IBM Transient Naming Server. When the component needs access to another business-integration-system component, it uses the naming service to determine the information it needs to locate and start interacting with that component. For example, when an adapter needs to communicate with InterChange Server, it obtains the location of InterChange Server through the Transient Naming Server.

**Note:** The IBM Transient Naming Server is part of the IBM Java ORB. Therefore, it is installed automatically on the InterChange Server machine as part of the WebSphere InterChange Server installation process.

This section provides the following information about managing the IBM Transient Naming Server:

- "Identifying registered components"
- "Using the Persistent Naming Server" on page 174

### Identifying registered components

The WebSphere InterChange Server product provides the CosNameServer\_Dump tool to list all valid ICS ORB objects currently registered with the IBM Transient Naming Server. This tool is located in the bin subdirectory of the product directory. You invoke it with the following command:

#### UNIX

```
CosNameServer_Dump.sh
```

## Windows

CosNameServer\_Dump.bat

### Using the Persistent Naming Server

When a component of the WebSphere business integration system starts, it registers itself with the IBM Transient Naming Server and its CORBA object is stored in the memory of the Transient Naming Server. However, if the Transient Naming Server goes down, its memory contents are lost. As a result, all components that had been registered with it would need to be rebooted so they could reregister with the naming service.

The Persistent Naming Server extends the capability of the IBM ORB Transient Naming Server so that the collection of CORBA objects that are registered with the Transient Naming Server are stored in a naming repository. The existence of the naming repository means that these CORBA references, rather than being only in the Transient Naming Server memory, are persistent; that is, they are available to other processes and ICS components in the event that the Transient Naming Server goes down. Other components do not need to shut down and restart in order to reregister with the naming service.

The default location of the naming repository is the following local file:  
ProductDir\CxCosNameRepos.ior.

You can change the location of this repository by setting the `CosNamingPersistenceFile` configuration parameter in the CORBA section of the ICS configuration file (`InterchangeSystem.cfg`).

By default, the Persistent Naming Server is enabled; that is, references to CORBA objects are maintained in the naming repository. However, for this naming server to run, you must explicitly start it with the `PersistentNameServer` startup file, located in the `bin` subdirectory of the product directory. This startup file takes the following steps:

1. Starts the IBM ORB Transient Naming Server
2. Starts the Persistent Naming Server to load the referenced CORBA objects into the naming repository.

As part of its startup process, InterChange Server updates the naming repository by copying the CORBA objects currently registered with the Transient Naming Server into the naming-repository file. When each adapter starts, it updates the naming repository with its information. If InterChange Server has not yet started when the adapter starts, the naming repository is updated whenever InterChange Server does start.

**Note:** If the Persistent Naming Server goes down, you can restart it with the `PersistentNameServer` startup script. However, you do not need to restart InterChange Server or any started adapters.

You can turn off the Persistent Naming Server by setting the `CosNamingPersistence` configuration parameter (also located in the CORBA section of the ICS configuration file) to `false`.



**High-Availability**

For the High-Availability (HA) environment, you *must* install the Persistent Naming Server as a Windows service. For more information, see the *System Installation Guide for Windows*.



---

## Appendix A. WebSphere MQ reference

This appendix describes some of the commands that are used at the DOS command prompt to administer WebSphere MQ. A complete description of the WebSphere MQ commands can be found in the WebSphere MQ online documentation.

UNIX command	Windows command	What the command does
start_mq	runmqsc or start_mq.bat	Starts the WebSphere MQ command interpreter. The remaining commands can be entered after this command is run.
end_mq	endmqm -i queue.manager	Ends WebSphere MQ queues immediately. Must be run before deleting queues.
clear_mq	clearorclear_mq.bat	Clears WebSphere MQ queues.
dltmqm queue.manager	dltmqm queue.manager	Deletes WebSphere MQ queues.
crtmqm queue.manager	crtmqm queue.manager	Creates an WebSphere MQ queue manager.
configure_mq <path to crossworlds_mq.tst>	configure_mq.bat <path to crossworlds_mq.tst>	Configures WebSphere MQ queue manager.
define	define	Defines WebSphere MQ queues.



---

## Appendix B. Requirements for restarting IBM WebSphere Business Integration system components

This appendix describes the restart requirements for administrative tasks. The restart requirements for development or implementation tasks are described in the appropriate development guide or in the *System Implementation Guide*.

Use the following tables to decide if it is necessary to restart individual IBM WebSphere InterChange Server Components, including InterChange Server. A Dynamic listing in the Restart Requirements column means the component does not require a restart.

- “InterChange Server restart requirements”
- “Collaboration template restart requirements”
- “Collaboration object restart requirements” on page 180
- “Connector restart requirements” on page 180
- “Business object restart requirements” on page 181
- “Map restart requirements” on page 182
- “Relationship restart requirements” on page 182

---

### InterChange Server restart requirements

Table 31 describes the restart requirements for InterChange Server.

**Note:** These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide*.

Table 31. Interchange Server restart requirements

Action	Restart requirement
Set Tracing options	Dynamic
Set Log/Trace file name	Dynamic
The max log/trace file size and number of archive files	Restart ICS
Change ICS name	Stop ICS, recreate queues with new names, then start ICS
Change ICS password	Dynamic
Change repository database passwords	Dynamic
Add new class libraries (*.jar)	Add the name to the startup file, then restart ICS
Change class libraries	Restart ICS

---

### Collaboration template restart requirements

Table 32 on page 180 describes the restart requirements for collaboration templates.

**Note:** These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide*.

Table 32. Collaboration template restart requirements

Action	Restart requirement
Recompile collaboration template without modifying the defined port names or adding or changing configuration property names.	Restart the collaboration object
Change port definition of the collaboration template.	Recompile the collaboration template; re-create the collaboration object
Change property names in the collaboration template.	Recompile the collaboration template; re-create the collaboration object

## Collaboration object restart requirements

Table 33 describes the restart requirements for collaboration objects.

**Note:** These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide*.

Table 33. Collaboration object restart requirements

Action	Restart requirement
Add a new collaboration object	Dynamic, but the user must bind ports to connectors, then start the collaboration object
Delete a collaboration object	Requires that the collaboration object be stopped
Change port bindings	Collaboration object must be stopped before bindings can be changed then restarted after the change is made
Change configuration property	Dynamic
Change reuse of collaboration object	Restart the collaboration object
Change system trace level	Dynamic
Change collaboration trace level	Dynamic
Add e-mail notification recipient	Dynamic
Change number of concurrent events	Dynamic
Change user defined properties	Dynamic

## Connector restart requirements

Table 34 describes the restart requirements for connectors.

**Note:** These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide*.

Table 34. Connector restart requirements

Action	Restart requirement
Add a new connector	Restart ICS
Delete a connector	Dependent collaboration objects must be deleted first
Change configuration property (other than those mentioned below)	Restart connector component. See property description in Connector Configurator to determine action
Change supported business objects properties	Restart connector
Remove supported business object	Restart connector
Change trace level of the connector agent	Dynamic
Change trace level of the connector controller	Dynamic
Change poll frequency	Dynamic
Change controller store and forward mode	Dynamic

Table 34. Connector restart requirements (continued)

Action	Restart requirement
Change transport protocol	Restart ICS
Change concurrent event triggered flows	Restart ICS
Change poll start time	Restart connector
Change poll end time	Restart connector agent
Create new connector	Does not require a restart of ICS
Recompile connector	Restart connector agent
New supported business objects	Restart connector agent
Bind maps to supported business object	Refresh connector screen in System Manager
Copy and paste connector	Restart ICS. Until ICS is restarted, the Associated Maps screen is blank because the information in this screen reflects the ICS's runtime only. Since this is a brand new connector with new supported BOs that are not yet reflected in ICS's runtime, the Associated Maps screen has no running BOs for binding maps.
Make changes to the values of the following standard properties:	From System Manager, right-click the server under Server Instances, then select Refresh. The Refresh option stops the connector, temporarily caches its transient state, removes the instance from memory and replaces it with the new instance, then starts the new instance.
<ul style="list-style-type: none"> <li>• DeliveryTransport</li> <li>• AllowAnonymousConnections</li> <li>• EbGateway</li> <li>• RemoteWebGatewayURL</li> <li>• ListenPort</li> <li>• CACertificateDirectory</li> <li>• ConcurrentEvents</li> <li>• JMSBrokerName</li> <li>• JMSFactoryClassName</li> <li>• JMSUserName</li> <li>• JMSPassword</li> <li>• NumberOfConnections</li> </ul>	

## Business object restart requirements

Table 35 describes the restart requirements for business objects.

**Note:** These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide*.

Table 35. Business object restart requirements

Action	Restart requirement
Add a business object	Dynamic
Delete a business object	Delete business object dependencies prior to deleting business object
Any change listed in this table.	If any tool, such as Map Designer or Process Designer, is connected to ICS when any change listed in this table is made, that tool must be disconnected from, then reconnected, to the server.
Change the application-specific text for attributes.	Restart connector agent. Restarting ICS is strongly recommended.
Change key attribute.	Restart connector agent.

Table 35. Business object restart requirements (continued)

Action	Restart requirement
Change default value.	If the connector uses the "UseDefaults" connector configuration property, the connector agent must be restarted. Restarting ICS is strongly recommended.
Change max length for an attribute.	Restart connector agent.
Change required field for an attribute.	Restart connector agent.
Change trace level.	Dynamic
Change attribute names.	Restart connector agent.
Change business-object structure (removing or adding attributes/sub-objects).	Stop ICS.  Clean queues.  Restart ICS.  The maps using the object must be updated and recompiled.

## Map restart requirements

Table 36 describes the restart requirements for maps.

**Note:** These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide*.

Table 36. Map restart requirements

Action	Restart requirement
Add or update map definitions	User must compile the map, which reloads the map definition, then restart the map
Update map properties	Restart the map
Delete a map	The map must be stopped before it can be deleted
Recompile the map	Dynamic Compilation of maps puts them in active state by default, regardless of map instance reuse option.
Change reuse of map instances	Restart the map
Change trace level and data validation level	Restart the map
Add imported class libraries in the map	If the class library is not already included in the startup script of the server, it must be included and ICS restarted.
Change imported class libraries	Restart ICS  If there are signature changes, the maps must also change and must be recompiled.

## Relationship restart requirements

Table 37 on page 183 describes the restart requirements for relationships.

**Note:** These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide*.



Table 37. Relationships restart requirements

Action	Restart requirement
Create or change a relationship.	Stop the relationship before saving the changes with “Create runtime schema” option, and then restart the relationship.



---

## Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Burlingame Laboratory Director  
IBM Burlingame Laboratory  
577 Airport Blvd., Suite 800

Burlingame, CA 94010  
U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples may include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

---

## Programming interface information

Programming interface information, if provided, is intended to help you create application software using this program.

General-use programming interfaces allow you to write application software that obtain the services of this program's tools.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

**Warning:** Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

---

## Trademarks and service marks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries, or both:

IBM  
the IBM logo  
AIX  
CrossWorlds  
DB2  
DB2 Universal Database  
Domino  
Lotus  
Lotus Notes  
MQIntegrator  
MQSeries  
Tivoli  
WebSphere

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

MMX, Pentium, and ProShare are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

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

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

The Monitor Definition Wizard and System Manager include software developed by the Eclipse Project (<http://www.eclipse.org>).



IBM WebSphere InterChange Server V4.2.2, IBM WebSphere Business Integration Toolset V4.2.2.



---

# Index

## A

- about logging 117
- about tracing 126
- archive tables, backing up 99
- automatic and remote restart of connectors 66

## B

- backing up
  - archive tables 99
  - collaboration class files 99
  - components 97
  - repository 98
- backup
  - planning 96
  - schedule 96
- bar display option, default monitors 8
- business object probe data file, deleting in System Monitor 22
- business objects
  - restart requirements 181
  - tracing 128

## C

- collaboration class files, backing up 99
- collaboration deadlocks
  - deleting 163
  - detecting group collaboration deadlocks 164
  - fixing 164
  - preventing 164
  - resolving 162
- collaboration objects
  - configuring concurrent event-triggered flow processing 74
  - configuring flow control 74
  - configuring run-time properties 72
  - configuring trace levels 135
  - log messages 118
  - management by SNMP agent 31
  - operating 69
  - restart requirements 180
  - scheduling 112
  - starting, stopping, and pausing 72
  - states 69
  - statistics 26
  - tracing 127
- collaboration templates, restart requirements 179
- community names
  - SNMP Agent Configuration Manager 36
- concurrent event-triggered flows, configuring for collaboration objects 74
- configuring
  - collaboration object trace level 135
  - collaboration tracing 135
  - connector agent logging destination 133
  - e-mail notification 123
  - flow control 54, 68, 74
  - InterChange Server logging and tracing parameters 132
  - InterChange Server logging and trace destinations 130
  - logging 129
  - map tracing 136

- configuring (*continued*)
  - persistent monitoring 39
  - stack trace 132
  - system-wide flow control 54
  - trace levels 134
  - tracing 129
- connecting to an InterChange Server instance 23
- connector agent
  - configuring logging destination 133
  - log messages 121
  - unknown status 150
- connectors
  - automatic and remote restart 66
  - commands for changing states 59
  - configuring flow control 68
  - configuring tracing 135
  - initializing 58
  - log messages 119
  - management by SNMP agent 32
  - operating 56
  - restart requirements 180
  - restarting 65
  - running, stopping, and pausing 58
  - scheduling 112
  - shutting down 64
  - starting manually 59
  - states 56
  - statistics 27
  - tracing 128
- creating schedules 113
- critical errors 148
- customizing the Object Request Broker (ORB) 170

## D

- database connection failures 151
- database connection pools, managing 166
- database, changing password 50, 51
- deleting schedules 115
- disabling schedules 115
- disconnecting from InterChange Server 85
- displaying schedules 115

## E

- e-mail notification
  - collaboration object level 125
  - configuring 123
  - connector level 126
  - system level 123
  - valid e-mail addresses 126
- enabling schedules 115
- event archives, maintaining 145

## F

- failed flows, resolving 151
- failed transactional collaborations 152
- failure recovery for service calls 146

- failures
  - database connection 151
  - flow 151
  - managing 145
- flow control
  - configuring for collaboration objects 74
  - configuring for connectors 68
  - configuring system-wide 54
- flow failures 151
- Flow Manager
  - finding unresolved flows 155
  - handling unresolved flows 160
  - managing queries 158
  - starting 153
  - viewing details of unresolved flows 158
- flow tracing 128

**H**

- High-Availability
  - maintaining a Windows HA system 168
  - managing 167
  - supported environments 167
- historical data, setting the frequency for capturing in System Monitor 20
- historical statistics, deleting in System Monitor 22

**I**

- IBM Transient Name Server 173
- initializing a connector 58
- installing
  - Object Request Broker (ORB) 170
  - SNMP Configuration Agent Manager 35
  - System Monitor 2
- InterChange Server
  - changing password 50, 51
  - configuring for flow monitoring 41
  - configuring logging and tracing parameters 132
  - connecting to an instance 23
  - customizing startup parameters 48
  - disconnecting from 85
  - graceful shutdown 50
  - how it works with WebSphere Business Integration Monitor 41
  - immediate shutdown 50
  - logs 120
  - management by SNMP agent 31
  - operating 48
  - password 50, 179
  - restart requirements 179
  - shutting down 49
  - starting 48
  - startup parameters 48
  - statistics 24
  - strategies for recovery 146
  - tracing services 127
- InterChange Server Component Management view of System Manager 22

**J**

- jobs, scheduling 110

**L**

- line display option, default monitors 8
- log file management 121
- log messages
  - viewing 137
- log viewer
  - using for tracing 137
- logging
  - about 117
  - collaboration object messages 118
  - configuring 129
  - configuring connector agent logging destination 133
  - configuring InterChange Server logging destination 130
  - connector agent messages 121
  - connector messages 119
  - High-Availability 121
  - message formats 119
  - system logs 120
- logging on to System Monitor 13
- long-lived business processes 75
- lost connection to application 150

**M**

- Management Information Base (MIB) 30
- managing
  - database connection pools 166
  - failures 145
  - High-Availability (HA) systems 167
  - log and trace files 121
  - Object Request Broker (ORB) 170
  - runtime properties 169
  - WebSphere MQ 165
  - WIP connections 145
- maps
  - configuring tracing 136
  - operating 76
  - restart requirements 182
  - starting and stopping 77
  - states 76
  - tracing 128
- meter display option, default monitors 9
- modifying schedules 114
- Monitor Definition Wizard, creating new monitors 9
- monitoring functions of InterChange Server Component Management view 22
- monitors
  - creating new 9
  - customizing the visual appearance in System Monitor 19
- monitors, default
  - bar display option 8
  - line display option 8
  - meterdisplay option 9
  - sample display options 6
  - stacked bar display option 9
  - table display option 7
  - table tree display option 7
  - using with System Monitor 2

**N**

- new in release 4.0.0 x
- new in release 4.0.1 ix
- new in release 4.1.0 ix
- new in release 4.1.1 ix
- new in release 4.2.0 viii



new in release 4.2.1 viii  
new in release 4.2.2 vii

## O

Object Request Broker (ORB)  
  changing location 172  
  customizing 170  
  installing 170  
  managing 170  
  using the IBM Transient Name Server 173  
opening relationships 85  
operating  
  collaboration objects 69  
  connectors 56  
  InterChange Server 48  
  maps 76  
  relationships 78  
overriding schedules 113

## P

password  
  changing for database 50, 51  
  changing for InterChange Server 50, 51  
  InterChange Server 50, 179  
persistent monitoring 39  
  accessing results 40  
  configuring 39

## R

reconfiguring SNMP agent 33  
refreshing System Manager 53  
relationship data  
  cleaning up participants 95  
  filtering displayed participants 94  
  printing relationship data 96  
  searching for participants 92  
Relationship Manager  
  starting 83  
  using 82  
relationships  
  copying participants 91  
  creating relationship instances 88  
  deactivating and activating participants 90  
  deleting relationship instances 90  
  loading and unloading business object files 91  
  opening 85  
  operating 78  
  restart requirements 182  
  retrieving relationship instances 86  
  starting and stopping 79, 80  
  states 78  
  table caching 80  
  working with 85  
  working with relationship data 92  
remote and automatic restart of connectors 66  
repos\_copy  
  arguments added in release 4.2 104  
  arguments removed in release 4.2 104  
  command 99  
  compiling components in the repository 107  
  deleting components from the repository 108  
  deploying a package to the repository 105  
  exporting components to a package 109

repos\_copy (*continued*)  
  locale 110  
  printing a list of components in the repository 109  
  printing the command 105  
  syntax 99  
  usage scenarios 104  
  validating the package 105  
  validating the repository 107  
Repository (InterChange Server)  
  loading 99  
repository, backing up 98  
required software, System Monitor 2  
resolving  
  collaboration deadlocks 162  
  failed flows 151  
restart requirements 179  
  business objects 181  
  collaboration objects 180  
  collaboration templates 179  
  connectors 180  
  InterChange Server 179  
  maps 182  
  relationships 182  
restarting connectors 65  
run-time properties, configuring for collaboration objects 72  
running, stopping, and pausing connectors 58

## S

sample display options, default monitors 6  
schedules  
  creating 113  
  deleting 115  
  displaying 115  
  enabling or disabling 115  
  modifying 114  
  overriding 113  
scheduling  
  collaboration objects 112  
  connectors 112  
  jobs 110  
server access entries, configuring in SNMP Agent  
  Configuration Manager 38  
server memory, controlling usage 143  
shutting down  
  connectors 64  
  InterChange Server 49  
  InterChange Server gracefully 50  
  InterChange Server immediately 50  
  System Manager 53  
SNMP agent  
  connecting to from the SNMP Agent Configuration  
  Manager 35  
  how SNMP agent and SNMP manager communicate 30  
  how SNMP works 29  
    community names 30  
    Management Information Base (MIB) 30  
    SNMP architecture 29  
  how to use 32  
  management of collaboration objects 31  
  management of connectors 32  
  management of InterChange Server 31  
  reconfiguring 33  
  starting and stopping 32  
  tracing 129  
  using 28  
  what SNMP manages 31

- SNMP Agent Configuration Manager 34
  - configuring community names 36
  - configuring server access entries 38
  - configuring trap forwarding 37
  - connecting to the SNMP agent 35
  - installing 35
  - starting 35
- SNMP manager, communication with SNMP agent 30
- stack trace, configuring 132
- stacked bar display option, default monitors 9
- starting
  - connectors manually 59
  - Flow Manager 153
  - InterChange Server 48
  - Relationship Manager 83
  - SNMP Agent Configuration Manager 35
  - System Manager 53
- starting and stopping
  - maps 77
  - relationships 79, 80
  - SNMP agent 32
- starting, stopping, and pausing
  - collaboration objects 72
- startup parameters
  - customizing for InterChange Server 48
  - InterChange Server 48
- state change log, deleting in System Monitor 22
- states
  - capturing state changes in System Monitor 22
  - collaboration objects 69
  - connectors 56
  - maps 76
  - relationships 78
  - relationships 78
- statistics
  - collaboration objects 26
  - connectors 27
  - InterChange Server 24
  - viewing and using from the InterChange Server
    - Component Management view 23
- System Manager
  - InterChange Server Component Management view 22
  - refreshing 53
  - shutting down 53
  - starting 53
  - updating components 53
  - using 52
- System Monitor
  - configuring existing views 18
  - creating new monitors 9
  - creating new views 17
  - deleting historical statistics 22
  - deleting the business object probe data file 22
  - deleting the state change log 22
  - deleting views 18
  - different approaches to using 2
  - fine-tuning 18
    - capturing state changes 22
    - customizing the visual appearance of monitors 19
    - resetting runtime statistic values 21
    - setting the default view 19
    - setting the frequency for capturing historical data 20
    - setting the refresh rate for runtime values 20
  - installing 2
  - logging on 13
  - required software 2
  - using default monitors 2

- System Monitor (*continued*)
  - using views 16
  - Web page contents 15

## T

- table caching, relationships 80
- table display option, default monitors 7
- table tree display option, default monitors 7
- trace file management 121
- tracing
  - about 126
  - business objects (flow tracing) 128
  - collaboration objects 127
  - configuring 129
  - configuring connector tracing 135
  - configuring InterChange Server tracing destination 130
  - configuring map tracing 136
  - configuring trace levels 134
  - configuring collaboration object trace levels 135
  - connectors 128
  - InterChange Server 127
  - maps 128
  - SNMP agent 129
  - using 136
  - using log viewer 137
  - Web gateways 129
- transactional collaborations, failed 152
- trap forwarding, configuring in SNMP Agent Configuration Manager 37
- typographic conventions v

## U

- unknown status, connector agent 150
- unresolved flows
  - finding 155
  - handling 160
  - viewing details 158
- updating components in System Manager 53

## V

- viewing log messages 137
- views
  - configuring existing in System Monitor 18
  - creating new in System Monitor 17
  - deleting in System Monitor 18
  - setting the default in System Monitor 19
  - using default in System Monitor 16
  - using in System Monitor 16

## W

- Web gateways, tracing 129
- Web page contents of System Monitor 15
- WebSphere Business Integration Monitor 40
  - administering 41
  - configuring InterChange Server for flow monitoring 41
  - how it works with InterChange Server 41
- WebSphere MQ
  - commands 177
  - errors 165
  - failure recovery and the queue manager 166
  - failures to create new instances of MQQueueManager 166

WebSphere MQ (*continued*)  
managing 165  
memory checker segmentation fault 166

WIP connections, managing 145  
working with relationships 85