




Installation and User's Guide



Installation and User's Guide

Note

Before using this information and the product it supports, read the information in Appendix C, "Notices," on page 45.

First Edition (April 2006)

This edition applies to Sparkler 1 of IBM Tivoli Composite Application Manager for SOA (product number 5724-M07 for the distributed version, and 5698-A77 for the Enterprise version) and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 2006. All rights reserved.

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

Contents

Tables	v
Chapter 1. Overview	1
Linux operating system support	2
What is included in Sparkler 1?	2
Related publications	3
Support information.	4
Chapter 2. Installing Sparkler 1	5
Installing Sparkler 1	5
Installing support for Tivoli Enterprise Portal Server	6
Installing support for Tivoli Enterprise Monitoring Server	6
Installing data collector support	7
Chapter 3. The JBoss data collector	9
Enabling the data collector	9
Disabling the data collector	11
Additional considerations	11
Chapter 4. The SCA data collector	13
Enabling the data collector	13
Using multiple data collectors in the same environment	15
Monitoring controls	15
Communication between SCA modules on different application servers	15
Logging and Filtering.	15
Disabling the data collector	16
Limitations	16
Chapter 5. The DataPower data collector	17
The DataPower data collector as a proxy	17
Planning for deployment	18
Aggregation	19
Deployment steps	19
Configuring the DataPower appliance for monitoring	20
Enabling the XML Management Interface	20
Checking additional optional settings	20
Configuring a user account on the DataPower appliance	20
Enabling the data collector	21
The DataPower configuration file	21
Running the KD4configDC command.	22
Specifying KD4configDC parameters	22
Creating node names in Tivoli Enterprise Portal	24
Examples of enabling the data collector.	25
Considerations for configuration of DataPower monitoring	26
Multiple connections to the same DataPower appliance	26
Restricting the list of domains being monitored	27
Separating and aggregating data with the <code>-displaygroup</code> parameter	27
Separating data by host name or display group	27
Separating data between domains by domain list	28
Aggregating data across multiple domains	28
Aggregating data across multiple domains under multiple nodes.	28
Aggregating data across all domains on a single host with multiple users	29
Aggregating data for a cluster of appliances	29

Additional considerations	30
Disabling the data collector	30
Communication between data collector and appliance	30
Optimizing performance	31
Limitations	31
Starting and stopping the data collector	31
Troubleshooting	32
Communication failures	32
Synchronizing time between computer systems	32
Password problems	32
Logging and Tracing	32
Appendix A. Messages	33
Appendix B. Obtaining IBM software support for IBM Tivoli Composite Application Manager for SOA	41
Searching knowledge bases	41
Searching the information center	41
Searching the Internet	41
Obtaining fixes	41
Receiving weekly support updates	42
Contacting IBM Software Support	42
Determining the business impact	43
Describing problems and gathering information	44
Submitting problems	44
Appendix C. Notices	45
Trademarks	47
Index	49

Tables

1. DataPower key and value pairs for KD4configDC 23

Chapter 1. Overview

IBM® Tivoli® Composite Application Manager for SOA provides the management function for managing the services and mediations in a service oriented architecture (SOA) running on Web application servers and the IBM WebSphere Enterprise Service Bus (WESB). For supported environments, IBM Tivoli Composite Application Manager for SOA monitors and performs simple control of message traffic between Web services in the SOA. IBM Tivoli Composite Application Manager for SOA relies on a distributed network of data collectors to feed data into Tivoli Management Services for analysis and management.

IBM Tivoli Composite Application Manager for SOA works in concert with other products to support the logical services management layer in a complete SOA management solution, and to manage the resources that support the services and the IBM WebSphere Enterprise Service Bus.

IBM Tivoli Composite Application Manager for SOA focuses on discovery of services in a wide range of application server runtime environments that are most common in early SOA adoption projects, including Microsoft .NET, BEA WebLogic Server, and IBM WebSphere Application Server. IBM Tivoli Composite Application Manager for SOA supports production IT environments, and status and situation generation.

IBM Tivoli Composite Application Manager for SOA is installed and operates within the management infrastructure of the Tivoli Management Services environment. The following list describes the primary components of IBM Tivoli Composite Application Manager for SOA:

- A *monitoring agent* that interacts with the managed application servers and infrastructure middleware to collect data, and store the data in a log file. The monitoring agent consists of the following sub-components:
 - An Intelligent Remote Agent (IRA)
 - One or more Data Collectors (DCs) that are installed locally on every application server runtime environment where Web services are to be managed. The data collector is the lowest-level component of the monitoring agent. This component is responsible for actually observing what is happening in the environment that it is designed to monitor. The data collected by the data collector is provided to the Data Collector Adapter through a log file on the monitored machine, for transmittal through the Tivoli Management Services infrastructure. Data collectors typically have access to the following types of information about the message traffic it monitors:
 - Source and destination (machine name, application server name, service name and operation name)
 - Whether the message is a request or a response
 - Interaction type (synchronous or asynchronous)
 - The association of a request to its response during an asynchronous interaction
 - The response time for the message
 - Whether or not the message generated a fault
 - Whether the message is a one-way or a two-way message
 - Optionally, the actual header and body content of the message itself
 - A Data Collector Adapter (DCA), used to communicate between the Data Collector and Intelligent Remote Agent. The Data Collector Adapter accesses

the data stored in the log file by the data collector, and stores data collector configuration data in a configuration file that is accessed by the data collectors on each managed system.

- A set of management data using logical table constructs.
- A set of queries and commands

IBM Tivoli Composite Application Manager for SOA can also process correlation information from the message traffic to create pattern and sequence views in the IBM Web Services Navigator, a separate tool provided with IBM Tivoli Composite Application Manager for SOA.

IBM Tivoli Composite Application Manager for SOA provides basic control of message traffic, including the logging of messages and filtering messages based on criteria that you can configure (machine name, service name, operation name, and IP address of the requester).

Linux operating system support

In addition to the operating system platforms supported by IBM Tivoli Composite Application Manager for SOA v6.0, additional support for Linux[®] operating systems is available through the installation of the IBM Tivoli Composite Application Manager for SOA v6.0 Linux Platforms Media Delivery Pack. This additional function provides support for running IBM Tivoli Composite Application Manager for SOA in all supported application server runtime environments on Linux operating systems. If you plan to install Sparkler 1 on a supported Linux operating system, you must first install this Media Delivery Pack on top of the base IBM Tivoli Composite Application Manager for SOA product.

See the IBM Passport Advantage[®] Web site for more information on obtaining and installing this Media Delivery Pack for Linux operating system support:

http://www-306.ibm.com/software/howtobuy/passportadvantage/pao_customers.htm

What is included in Sparkler 1?

Version 6.0 of IBM Tivoli Composite Application Manager for SOA provides support for monitoring and simple control of Web services running in IBM WebSphere Application Server, BEA WebLogic Server and Microsoft .NET. This support includes a set of data collectors within the OMEGAMON[®] infrastructure, taking advantage of functions available in the Tivoli Enterprise Portal. These data collectors also operate in the IBM Tivoli Monitoring v6.1 infrastructure.

A *Sparkler* is an IBM delivery mechanism used to include additional features in a current product release. This document describes the installation and operation of the additional features included in Sparkler 1 for IBM Tivoli Composite Application Manager for SOA v6.0. Sparkler features are provided with limited function.

Sparkler 1 broadens the coverage of supported application server runtime environments, providing the following additional features:

- Support for JAX-RPC Web services running in the JBoss 4.0.3 application server runtime environment.
- Monitoring support for Web services flows through a DataPower appliance acting as a proxy between the Web service client and server.

- Instrumentation and monitoring support for the Services Component Architecture (SCA) infrastructure of IBM WebSphere Process Server v6.0.1 and IBM WebSphere Enterprise Service Bus v6.0.1.

These new application server runtime environments provide similar function as those supported by IBM Tivoli Composite Application Manager for SOA v6.0, and are displayed as additional Environment types in Tivoli Enterprise Portal workspaces and views, with the following limitations:

- The SCA data collector does not provide the same level of functionality as other data collectors.
- SCA data is displayed in the same set of existing workspaces and views as the IBM WebSphere Application Server data collector.
- The DataPower data collector does not provide any filtering capabilities.
- There is no support in the IBM Web Services Navigator for the DataPower and SCA data collectors.

The rest of this document describes the procedure for installing Sparkler 1 and provides more information on configuring and enabling the data collectors to run in these new application server runtime environments.

Related publications

This section lists publications in the IBM Tivoli Composite Application Manager for SOA library. Refer to this documentation for more information about the base product for this Sparkler 1, and for additional information on how to access IBM Tivoli publications online and how to order IBM Tivoli publications.

The IBM Tivoli Composite Application Manager for SOA library contains the following publications:

- *IBM Tivoli Composite Application Manager for SOA Release Notes*, GI11-4096
Contains late-breaking information about IBM Tivoli Composite Application Manager for SOA product limitations and workarounds, and pointers to other documentation to help you begin installing and using the product.
- *IBM Tivoli Composite Application Manager for SOA Installation and User's Guide*, GC32-9492

Provides an overview of the Tivoli Management Services environment and the planning information and procedures you need to install and configure the IBM Tivoli Composite Application Manager for SOA monitoring agent. This guide also provides use information on monitoring and managing resources with IBM Tivoli Composite Application Manager for SOA. This guide also provides troubleshooting information about the various components of IBM Tivoli Composite Application Manager for SOA, as well as information about log files and informational, warning, and error messages.

- *Configuring IBM Tivoli Composite Application Manager for SOA on z/OS®*, SC32-9493

Provides information about configuring IBM Tivoli Composite Application Manager for SOA to operate in the z/OS operating system environment. Before using this publication, you must complete the installation procedures as documented in the *Program Directory for IBM Tivoli Composite Application Manager for SOA, V6.0, Program Number 5698-A77, for Use with z/OS*, GI11-4100. This publication is written for system administrators and others who are responsible for installing and configuring IBM Tivoli Composite Application Manager for SOA in the z/OS environment.

- *Installing and Troubleshooting IBM Web Services Navigator*, GC32-9494
Provides information about installing and using an Eclipse based plug-in for extracting Web services information that has been collected by IBM Tivoli Composite Application Manager for SOA data collectors and stored in a database, and displaying the data in several views to assist a Web services architect in visualizing relationships between services.

Support information

Appendix B, “Obtaining IBM software support for IBM Tivoli Composite Application Manager for SOA,” on page 41 describes the following options for obtaining support for IBM products:

- “Searching knowledge bases” on page 41
- “Obtaining fixes” on page 41
- “Contacting IBM Software Support” on page 42

Chapter 2. Installing Sparkler 1

The installation of Sparkler 1 involves downloading the required zipped packages from the product support Web site, and unpacking them into existing IBM Tivoli Composite Application Manager for SOA product directories, then configuring the Tivoli Enterprise Portal Server, Tivoli Enterprise Monitoring Server, and data collectors (for each supported platform) as needed.

This installation assumes that you have already successfully installed IBM Tivoli Composite Application Manager for SOA v6.0 in your Tivoli Management Services environment. See “Related publications” on page 3 for installation and configuration procedures.

Installing on Linux operating systems: If you are installing this Sparkler 1 on a supported Linux operating system, this installation also assumes that you have already installed the IBM Tivoli Composite Application Manager for SOA v6.0 Linux Platforms Media Delivery Pack for this support before installing this Sparkler. See “Linux operating system support” on page 2 for more information on this prerequisite.

Do you have a distributed installation? If so, install Sparkler 1 first on the computer system where Tivoli Enterprise Portal Server is installed, followed by the Tivoli Enterprise Monitoring Server system, and then the systems where data collectors are located.

Note that the Tivoli Enterprise Monitoring Server does not need to be reseeded after installing Sparkler 1.

Are other data collectors already enabled? The computer system where you install Sparkler 1 might already have other IBM Tivoli Composite Application Manager for SOA data collectors enabled and running. Any combination of supported IBM Tivoli Composite Application Manager for SOA data collectors can coexist on a single system. You can enable as many instances of each data collector as you choose.

Re-enable existing data collectors after installing: Due to shared code between the data collectors provided with this Sparkler 1 and other data collectors provided with IBM Tivoli Composite Application Manager for SOA (IBM WebSphere Application Server and BEA WebLogic Server only), after installing this Sparkler, you must run the KD4configDC command to disable and re-enable these types of existing data collectors that are deployed on the same system, to ensure that all data collectors are running at the same product level. See the documentation library for IBM Tivoli Composite Application Manager for SOA for details on disabling and re-enabling these other data collectors.

See the *IBM Tivoli Composite Application Manager for SOA Release Notes* and the *ReadMe* for the Linux patch for information on supported operating system platforms and other hardware and software requirements.

Installing Sparkler 1

Complete the following steps to install this Sparkler into your existing IBM Tivoli Composite Application Manager for SOA product environment:

1. This procedure assumes that you have already installed the base IBM Tivoli Composite Application Manager for SOA product in your Tivoli Management Services environment. See the publications provided with IBM Tivoli Composite Application Manager for SOA for more information.
2. If you are installing on a Linux operating system supported by the IBM Tivoli Composite Application Manager for SOA v6.0 Linux Platforms Media Delivery Pack, it is assumed that you have already applied this delivery pack to the base IBM Tivoli Composite Application Manager for SOA product. See the documentation provided with the delivery pack for more information.
3. The installation of Sparkler 1 involves downloading and unpacking a set of compressed files into the existing Tivoli Management Services environment. Make sure you follow your local enterprise procedures to backup the Tivoli Management Services environment prior to installing Sparkler 1 in case you need to restore your environment to the current state.
4. **Is a BEA WebLogic Server running?:** If you already have BEA WebLogic Server running on the same target computer system where you are installing this Sparkler, stop the BEA WebLogic Server before installing Sparkler 1.
5. If you are installing support for Tivoli Enterprise Portal Server, Tivoli Enterprise Monitoring Server, or data collector support on each computer system where Web services are to be monitored, see the sections that follow.

Installing support for Tivoli Enterprise Portal Server

To install support for Tivoli Enterprise Portal Server, complete the following steps:

1. Stop the Tivoli Enterprise Portal Server.
2. Download and unpack the appropriate compressed file package as follows:
 - For supported Windows[®] operating systems:
 - a. Locate and download the **6.0.0-TIV-ITCAMSOAS1_TEPS-WIN-FT001.zip** compressed file package to a temporary directory on the Tivoli Enterprise Portal Server computer system.
 - b. Unpack the files into the *<TMS_DIR>* directory, where *<TMS_DIR>* is the directory where Tivoli Management Services is installed (for example, C:\ibm\itm or C:\Candle on Windows). The files will be unpacked into their correct location in the \CNPS subdirectory.
 - For supported Linux operating systems:
 - a. Locate and download the **6.0.0-TIV-ITCAMSOAS1_TEPS-UNIX-FT001.tar.gz** compressed file package to a temporary directory *<TEMP_DIR>* on the Tivoli Enterprise Portal Server computer system.
 - b. Unpack the compressed file package with the following command:


```
gzip -d 6.0.0-TIV-ITCAMSOAS1_TEPS-UNIX-FT001.tar.gz
```
 - c. Navigate to the \$CANDLE_HOME/li6243 directory where Tivoli Enterprise Portal Server is installed.
 - d. Verify that there is a subdirectory named /cq in this directory.
 - e. Run the following command:


```
tar -xvf <TEMP_DIR>/6.0.0-TIV-ITCAMSOAS1_TEPS-UNIX-FT001.tar
```
3. Restart the Tivoli Enterprise Portal Server.

Installing support for Tivoli Enterprise Monitoring Server

To install support for Tivoli Enterprise Monitoring Server, complete the following steps:

1. Stop the Tivoli Enterprise Monitoring Server.

2. Download and unpack the appropriate compressed file package as follows:
 - For supported Windows operating systems:
 - a. Locate and download the **6.0.0-TIV-ITCAMSOAS1_TEMS_CMS-WIN-FT001.zip** compressed file package to a temporary directory on the Tivoli Enterprise Monitoring Server computer system.
 - b. Unpack the files into the *<TMS_DIR>* directory, where *<TMS_DIR>* is the directory where Tivoli Management Services is installed (for example, C:\ibm\itm or C:\Candle). The files will be unpacked into their correct location in the \CMS subdirectory.
 - For supported AIX, Solaris, or Linux operating systems:
 - a. Locate and download the **6.0.0-TIV-ITCAMSOAS1_TEMS_CMS-UNIX-FT001.tar.gz** compressed file package to a temporary directory *<TEMP_DIR>* on the Tivoli Enterprise Monitoring Server computer system.
 - b. Unpack the compressed file package with the following command:


```
gzip -d 6.0.0-TIV-ITCAMSOAS1_TEMS_CMS-UNIX-FT001.tar.gz
```
 - c. Navigate to the \$CANDLE_HOME/tables/*<TEMS_Server>* directory, where *<TEMS_Server>* is the name of the Tivoli Enterprise Monitoring Server.
 - d. Verify that there are subdirectories named /ATTRIB and /RKDSCATL under this subdirectory.
 - e. Run the following command:


```
tar -xvf <TEMP_DIR>/6.0.0-TIV-ITCAMSOAS1_TEMS_CMS-UNIX-FT001.tar
```
3. Restart the Tivoli Enterprise Monitoring Server.

Installing data collector support

To install data collector support on each computer system where services are to be monitored, or where the DataPower data collector proxy will be running (see Chapter 5, “The DataPower data collector,” on page 17 for more information), complete the following steps:

1. Stop the IBM Tivoli Composite Application Manager for SOA service.
2. Download and unpack the appropriate compressed file package as follows:
 - For supported Windows operating systems, locate and download the **6.0.0-TIV-ITCAMSOAS1_DCA-WIN-FT001.zip** compressed file package to a temporary directory, and then unpack the files into the *<TMS_DIR>* directory. The files will be unpacked into their correct location in the \CMA subdirectory.
 - For supported AIX[®] operating systems, locate and download the **6.0.0-TIV-ITCAMSOAS1_DCA-AIX433-FT001.tar.gz** compressed file package to a temporary directory, and then unpack the files into the *<TMS_DIR>* directory. The files will be unpacked into their correct location in the /aix433 subdirectory.
 - For supported Solaris operating systems, locate and download the **6.0.0-TIV-ITCAMSOAS1_DCA-SOL283-FT001.tar.gz** compressed file package to a temporary directory, and then unpack the files into the *<TMS_DIR>* directory. The files will be unpacked into their correct location in the /sol283 subdirectory.
 - For supported Linux (on Intel[®]) operating systems, locate and download the **6.0.0-TIV-ITCAMSOAS1_DCA-LI6243-FT001.tar.gz** compressed file package to a temporary directory, and then unpack the files into the *<TMS_DIR>* directory. The files will be unpacked into their correct location in the /li6243 subdirectory.

- For supported Linux (on z/OS) operating systems, locate and download the **6.0.0-TIV-ITCAMSOAS1_DCA-LS3243-FT001.tar.gz** compressed file package to a temporary directory, and then unpack the files into the *<TMS_DIR>* directory. The files will be unpacked into their correct location in the /ls3243 subdirectory.
3. If there are existing IBM Tivoli Composite Application Manager for SOA data collectors that are enabled on these computer systems, re-enable them by completing the following basic steps:
 - a. Stop the applications servers that currently have an IBM Tivoli Composite Application Manager for SOA data collector enabled.
 - b. Re-enable the data collector by running the KD4configDC command (see the IBM Tivoli Composite Application Manager for SOA documentation for details).
 - c. Restart the application servers.
 4. Restart the IBM Tivoli Composite Application Manager for SOA service.

The chapters that follow describe the new data collector support provided with Sparkler 1.

Chapter 3. The JBoss data collector

This chapter describes the support for JAX-RPC Web services running in the JBoss Java™ 2 Enterprise Edition (J2EE) version 4.0.3 application server runtime environment.

Enabling the data collector

To enable the JBoss data collector for data collection, complete the following steps:

1. You should have already completed the installation procedure for installing Sparkler 1 in the existing IBM Tivoli Composite Application Manager for SOA environment on the computer where your JBoss application server is installed. See “Installing Sparkler 1” on page 5 for details.
2. Optionally stop the JBoss application server before enabling the data collector. You do not have to stop the application server before running the KD4configDC command, but the data collector does not begin to collect data until after the JBoss application server is stopped and restarted. You might prefer to stop and restart the JBoss application server during off shift hours. Refer to your JBoss documentation for the specific procedure to stop the JBoss application server.
3. Run the KD4configDC command to enable the data collector. When there are multiple server instances within the same JBoss application server runtime environment, you must run the KD4configDC command for each instance. Run the KD4configDC command as follows:

- a. Depending on your operating system, navigate to one of the following locations:

- For supported Windows operating systems, navigate to %CANDLE_HOME%\CMA\KD4\bin.
- For supported AIX operating systems, navigate to \$CANDLEHOME/aix433/d4/KD4/bin.
- For supported Solaris operating systems, navigate to \$CANDLEHOME/sol283/d4/KD4/bin.
- For supported Linux (on Intel) operating systems, navigate to \$CANDLEHOME/li6243/d4/KD4/bin.
- For supported Linux (on z/OS) operating systems, navigate to \$CANDLEHOME/ls3243/d4/KD4/bin.

- b. Run the following command:

- For supported Windows operating systems:
KD4configDC.bat -enable -env 4 <server configuration> <JBoss_Home>
- For supported AIX, Solaris, and Linux operating systems:
./KD4configDC.sh -enable -env 4 <server configuration> <JBoss_Home>

In these commands, the following parameters are specified:

- <server configuration> is the type of JBoss application server to be configured. Valid values are:

default

This configuration type contains everything needed to run a standalone J2EE server.

all

This configuration type starts all available services, including the RMI/IIOP and clustering services, and the Web services deployer, which are not loaded in the *default* configuration

A third possible configuration type, **minimal**, is not supported.

You can also create your own server configuration type by modifying the *default* or *all* types, or mixing and matching capabilities from either or both, and saving it as a new server configuration type name. You can then specify that new server configuration type in the KD4configDC command.

- *<JBoss_Home>* is the directory path name where the JBoss application server is installed, for example, *C:\JBoss* on a Windows operating system (note that if this path name contains a blank space, the entire path is surrounded with quotation marks). If you have this directory path defined in the environment variable, *JBOSS_HOME*, then this parameter is optional.

Examples:

- `KD4configDC.bat -enable -env 4 default`

In this example, the **default** server configuration is used, and the *JBOSS_HOME* environment variable is used to specify the JBoss base installation directory. The command completes with a return code of 0.

- `KD4configDC.bat -enable -env 4 all C:\JBoss`

In this example, the **all** server configuration type is specified to enable monitoring for the JBoss configuration that has all services started. The JBoss base installation directory is explicitly specified, and because there are no blank spaces in the path, the argument can be specified without surrounding quotation marks.

- `KD4configDC.bat -enable -env 4 myConfig`

In this example, the user-customized server configuration type *myConfig* is specified, and the *JBOSS_HOME* environment variable is used by default to specify the JBoss base installation directory.

- `KD4configDC.bat -enable -env 4 default "C:\App Servers\JBoss"`

In this example, the JBoss base installation directory is specified explicitly, and because there is a blank space in the path, the argument is surrounded by quotation marks.

- `./KD4configDC.sh -enable -env 4 all /opt/IBM/JBoss`

This example is for running KD4configDC on a supported AIX, Solaris, or Linux operating system.

You can issue the following command to view the online help for KD4configDC:

- For supported Windows operating systems:

```
KD4configDC -h
```

- For supported AIX, Solaris, or Linux operating systems:

```
./KD4configDC.sh -h
```

4. Restart the JBoss application server (refer to your JBoss documentation for the specific procedure). The data collector becomes active after the JBoss application server is restarted.

The JBoss data collector only needs to be enabled once in the JBoss application server runtime environment. If new applications are added to the environment after the data collector has been enabled, the data collector begins monitoring the new applications automatically.

Disabling the data collector

You can disable the JBoss data collector using the `KD4configDC` command as follows:

1. Optionally stop the JBoss application server.
2. Depending on your operating system, navigate to one of the following locations:
 - For supported Windows operating systems, navigate to `%CANDLE_HOME%\CMA\KD4\bin`.
 - For supported AIX operating systems, navigate to `$CANDLEHOME/aix433/d4/KD4/bin`.
 - For supported Solaris operating systems, navigate to `$CANDLEHOME/sol283/d4/KD4/bin`.
 - For supported Linux operating systems, navigate to `$CANDLEHOME/li6243/d4/KD4/bin`.
3. Run the following command:
 - For supported Windows operating systems:
`KD4configDC.bat -disable -env 4 <server configuration> <JBoss_Home>`
 - For supported AIX, Solaris, and Linux operating systems:
`./KD4configDC.sh -disable -env 4 <server configuration> <JBoss_Home>`
4. After the command completes, restart the JBoss application server.

The JBoss data collector continues to collect data until after the JBoss application server is restarted.

Additional considerations

As you enable and use the JBoss data collector, keep in mind the following additional considerations:

- If the JBoss application server is stopped, you can successfully enable or disable the data collector.
- If you attempt to enable or disable the data collector multiple times in a row (for example, an enable followed by another enable), only the first invocation takes effect.

See the documentation library and online help information provided with IBM Tivoli Composite Application Manager for SOA for information on starting and stopping data collectors, configuring monitoring and filtering settings, understanding workspaces, views, Take Action commands, and other features that help you monitor and manage your Web services in the JBoss application server runtime environment.

Chapter 4. The SCA data collector

This chapter describes the support for instrumenting and monitoring the Service Component Architecture (SCA) infrastructure of IBM WebSphere Process Server and IBM WebSphere Enterprise Service Bus. Note that this Sparkler does not provide full data collector support for the SCA environment.

IBM WebSphere Process Server and IBM WebSphere Enterprise Service Bus introduce a new way to model services in an SOA, called the *Service Component Architecture* (SCA). SCA was designed to separate business logic from its implementation, so that you can focus on assembling an integrated application without knowing implementation details.

The Service Component Architecture is based on *SCA modules* and *SCA components*. An SCA module is made up of multiple SCA components. In IBM Tivoli Composite Application Manager for SOA, SCA components are treated as *services*.

With additional support for the IBM WebSphere Enterprise Service Bus and IBM WebSphere Process Server in Sparkler 1, IBM Tivoli Composite Application Manager for SOA discovers information about when a message flows between SCA components.

The SCA data collector is installed and configured once in the IBM WebSphere Process Server or IBM WebSphere Enterprise Service Bus environment. The data collector does not have to be configured for each SCA application. When new SCA applications are added to the monitored environment, they are automatically monitored as well.

The SCA data collector supports both synchronous and asynchronous interactions flowing through the SCA application server runtime environment. Note that some asynchronous interactions between SCA components might result in requests or responses being reported multiple times or not at all. This is caused by thread switching in the SCA application server runtime environment and limitations of the SCA data collector in tracking flows across these thread switches. Monitoring of asynchronous interactions is provided in this Sparkler to give you a more complete picture of the services that exist in the environment, however you should not expect the metrics collected for asynchronous interactions to be precise.

Refer to your IBM WebSphere Integration Developer documentation for more information about these synchronous and asynchronous interactions.

Enabling the data collector

To enable the SCA data collector for data collection, complete the following steps:

1. You should have already completed the installation procedure for installing Sparkler 1 in the existing IBM Tivoli Composite Application Manager for SOA environment on the computer where the IBM WebSphere Enterprise Service Bus and IBM WebSphere Process Server application server runtime environments are installed. See “Installing Sparkler 1” on page 5 for details.
2. Optionally stop the IBM WebSphere Enterprise Service Bus or IBM WebSphere Process Server application server before enabling the data collector. You do not have to stop the application server before running the KD4configDC command, but the data collector does not begin to collect data until after the application server is stopped and restarted. You might prefer to stop and restart the

application server during off shift hours. Refer to your WebSphere® documentation for the specific procedure to stop the application server.

3. Run the KD4configDC command as follows:
 - a. Depending on your operating system, navigate to one of the following locations:
 - For supported Windows operating systems, navigate to %CANDLE_HOME%\CMA\KD4\bin.
 - For supported AIX operating systems, navigate to \$CANDLEHOME/aix433/d4/KD4/bin.
 - For supported Solaris operating systems, navigate to \$CANDLEHOME/sol283/d4/KD4/bin.
 - For supported Linux (on Intel) operating systems, navigate to \$CANDLEHOME/li6243/d4/KD4/bin.
 - For supported Linux (on z/OS) operating systems, navigate to \$CANDLEHOME/ls3243/d4/KD4/bin.

b. Run the following command:

- For supported Windows operating systems:
KD4configDC.bat -enable -env 9 <WESB_HOME> | <WPS_HOME>
- For supported AIX, Solaris, and Linux operating systems:
./KD4configDC.sh -enable -env 9 <WESB_HOME> | <WPS_HOME>

In these commands, the following parameters are specified:

- <WESB_HOME> is the directory path name where IBM WebSphere Enterprise Service Bus is installed.
- <WPS_HOME> is the directory path name where IBM WebSphere Process Server is installed.

If the directory path name contains a blank space, the entire path must be surrounded with quotation marks. If you have this directory path defined in an environment variable, WBI_HOME, then this parameter is optional.

Examples:

- KD4configDC.bat -enable -env 9 %WBI_HOME%
In this example, the %WBI_HOME% environment variable is used to specify the IBM WebSphere Enterprise Service Bus base installation directory. The command completes with a return code of 0.
- KD4configDC.bat -enable -env 9 C:\WPS
In this example, the IBM WebSphere Process Server base installation directory is explicitly specified, and because there are no blank spaces in the path, the argument can be specified without surrounding quotation marks.
- KD4configDC.bat -enable -env 9 "C:\Program Files\WPS"
In this example, the IBM WebSphere Process Server base installation directory is specified explicitly, and because there is a blank space in the path, the argument is surrounded by quotation marks.
- ./KD4configDC.sh -enable -env 9 /opt/IBM/WPS
This example is for running KD4configDC on a supported AIX, Solaris, or Linux operating system.

You can issue the following command to view the online help for KD4configDC:

- For supported Windows operating systems:
KD4configDC -h
- For supported AIX, Solaris, or Linux operating systems:
./KD4configDC.sh -h

- Restart the IBM WebSphere Enterprise Service Bus or IBM WebSphere Process Server application server (refer to your documentation for the specific procedure). The data collector becomes active after the application server is restarted.

Using multiple data collectors in the same environment

The IBM WebSphere Application Server data collector provided with IBM Tivoli Composite Application Manager for SOA v6.0 can also be used in the IBM WebSphere Enterprise Service Bus and IBM WebSphere Process Server environments to monitor Web service flows, similar to basic IBM WebSphere Application Server environments.

When used together, the SCA data collector and the existing IBM WebSphere Application Server data collector provide monitoring for both SCA flows and Web services flows, respectively, in the IBM WebSphere Enterprise Service Bus and IBM WebSphere Process Server application server runtime environments. IBM recommends that you install and enable both of these data collectors in your WebSphere environment. Each of these data collectors can be independently enabled and disabled.

Because the SCA and IBM WebSphere Application Server data collectors run in the same application server runtime environment, they share the same configuration settings, they share a common set of log files (metric, operational, and trace), and they appear as one application server node in the Tivoli Enterprise Portal.

Monitoring controls

Using the Tivoli Enterprise Portal controls to turn monitoring on or off for the WebSphere application server runtime environment results in both the IBM WebSphere Application Server and SCA data collectors being turned on or off. The monitor controls defined for the WebSphere application server runtime environment apply to both the IBM WebSphere Application Server and SCA data collectors (though for SCA, when monitoring is enabled, the only valid message content level option is *none*). Any changes to the configuration parameters in the Tivoli Enterprise Portal views affect both data collectors.

Communication between SCA modules on different application servers

The SCA data collector monitors flows between SCA components within a single SCA module. For flows between SCA modules on different application servers, IBM Tivoli Composite Application Manager for SOA can only monitor flows that use Web services bindings. These flows are monitored only if the IBM WebSphere Application Server data collector for Web service flows is also enabled. Other types of bindings between SCA modules are not monitored by IBM Tivoli Composite Application Manager for SOA.

IBM recommends that you enable both the SCA and IBM WebSphere Application Server data collectors in your IBM WebSphere Enterprise Service Bus or IBM WebSphere Process Server application server runtime environments.

Logging and Filtering

The SCA data collector supports the monitor controls to define which SCA services and operations to monitor, but ignores the message content logging setting (*none*, *headers*, *body*, *full*) and never records message content.

The filter controls (rejection of messages) are not supported by the SCA data collector. This function is supported by IBM WebSphere Enterprise Service Bus and IBM WebSphere Process Server mediations to manage the rejection of messages. To record message content, use the log mediation function of IBM WebSphere Enterprise Service Bus or IBM WebSphere Process Server.

Disabling the data collector

You can disable the SCA data collector using the KD4configDC command. Optionally stop the WebSphere application server and run the KD4configDC command as follows:

1. Depending on your operating system, navigate to one of the following locations:
 - For supported Windows operating systems, navigate to %CANDLE_HOME%\CMA\KD4\bin.
 - For supported AIX operating systems, navigate to \$CANDLEHOME/aix433/d4/KD4/bin.
 - For supported Solaris operating systems, navigate to \$CANDLEHOME/sol283/d4/KD4/bin.
 - For supported Linux operating systems, navigate to \$CANDLEHOME/li6243/d4/KD4/bin.
2. Run the following command:
 - For supported Windows operating systems:
`KD4configDC.bat -disable -env 9 <WBI_HOME> | <WPS_HOME>`
 - For supported AIX, Solaris, and Linux operating systems:
`./KD4configDC.sh -disable -env 9 <WBI_HOME> | <WPS_HOME>`

After the command completes, restart the application server. The data collector continues to collect data until the application server is restarted.

Limitations

You should be aware of the following limitations for this Sparkler 1 version of the SCA data collector:

- The SCA data collector only supports monitoring, not filtering. You can define monitoring controls to limit which SCA services or operations are monitored, but any filter controls defined to reject messages are ignored.
- Monitor control for the message content parameter (*none, headers, body, full*) only supports the *none* value. All other values are ignored and treated as *none*.
- The message length for SCA interactions is always reported as zero.
- SCA data collector information might not render properly in the IBM Web Services Navigator topology, sequence diagram and patterns views.
- Some asynchronous interactions between SCA components might result in requests or responses being reported multiple times or not at all. This is caused by thread switching in the SCA application server runtime environment and limitations of the SCA data collector in tracking flows across these thread switches. Monitoring of asynchronous interactions is provided in this Sparkler to give you a more complete picture of the services that exist in the environment, however you should not expect the metrics collected for asynchronous interactions to be precise.

Chapter 5. The DataPower data collector

This chapter describes the support for monitoring of Web services flows through a DataPower appliance acting as a proxy between the Web services client and server.

IBM recently acquired DataPower, a manufacturer of XML and Web services acceleration and security *appliances*. Using a DataPower appliance, you can improve the security and performance of Web services by offloading functions from the application server that is hosting the target Web service to a DataPower appliance. Typical functions that are off-loaded include authentication and authorization, XML schema validation, and Web services encryption and decryption.

Sparkler 1 for IBM Tivoli Composite Application Manager for SOA v6.0 provides a new DataPower data collector that monitors Web services flows through a DataPower appliance and provides similar services management and availability information that IBM Tivoli Composite Application Manager for SOA currently provides for application server runtime environments. This information is displayed in the Tivoli Enterprise Portal using the usual predefined or user-defined workspaces and views.

Upgrade your firmware: Before using the DataPower data collector, you must upgrade the firmware on DataPower appliances that you want to monitor, to include the necessary monitoring and data transformation capabilities. Be sure to upgrade your firmware to version **3.5.0.5** or later.

The DataPower data collector as a proxy

Data collectors provided with IBM Tivoli Composite Application Manager for SOA are usually installed directly into the application server runtime environment hosting the services being monitored. The DataPower appliance, however, does not support the installation of additional software, such as a data collector. Instead, the DataPower appliance provides a communication mechanism that allows external software applications to receive data from its internal transaction log.

The DataPower data collector is installed on a separate machine and uses this communication mechanism to retrieve monitoring data about Web services requests flowing through one or more DataPower appliances, and converts the data into a format that IBM Tivoli Composite Application Manager for SOA can process. In this way, the DataPower data collector acts as a *proxy* between the DataPower appliances and the IBM Tivoli Composite Application Manager for SOA intelligent remote agent (IRA). The DataPower data collector can be installed on a dedicated machine, or it can run on a machine that is also running data collectors for other application server runtime environments. Note that IBM supports only one instance of the DataPower data collector running on any given machine.

When the DataPower data collector starts, it subscribes to each monitored DataPower appliance and then polls (at an interval that you can configure) the appliance for monitoring data. The data retrieved from the DataPower appliance is written to metric log files in the format used by IBM Tivoli Composite Application Manager for SOA. When this data is later displayed in the Tivoli Enterprise Portal, nodes are displayed in the Tivoli Enterprise Portal Navigation view that represent the DataPower appliances being monitored. You can select workspaces under these nodes and view the service management data for the Web services requests flowing through the monitored DataPower appliances.

The DataPower data collector can subscribe to multiple DataPower appliances, and retrieve and manage data from multiple domains. This data can then be separated by DataPower domain or aggregated across multiple domains and appliances, depending on how you configure the data collector. The DataPower data collector uses a configuration file that contains information about which DataPower appliances are being monitored and information needed to establish communication with each monitored appliance.

Planning for deployment

DataPower Web services proxies are defined within *application domains*, and DataPower users can be restricted to access some or all domains. When configuring the DataPower data collector, you must understand how the domains and users are defined on the monitored DataPower appliances, to ensure that the data collector uses valid authentication credentials. This refers to user IDs and passwords that have access to the DataPower domains containing the Web services proxies to be monitored. In addition, you must decide how you want to aggregate or separate the data collected from those domains for display in the Tivoli Enterprise Portal.

You can use DataPower appliances in several typical configurations:

- **Single appliance, single domain:** The data collector monitors a single DataPower appliance, with all of the monitored resources defined in a single domain on the appliance.
- **Single appliance, multiple domains:** The data collector monitors a single DataPower appliance, but that appliance has monitored resources that are defined in more than one domain on the appliance.
- **Multiple appliances with different configurations:** The data collector monitors multiple DataPower appliances, and each appliance has a different configuration of resources to be monitored. Each appliance is configured for a particular job, with no intention of load-balancing or fail-over between appliances.
- **Multiple appliances with identical configurations:** The data collector monitors multiple DataPower appliances, and all of the appliances have an identical configuration of resources being monitored. All of the appliances are configured for the same job, taking advantage of load-balancing and fail-over capabilities between appliances.

Given these typical configurations, the DataPower data collector provides a great deal of flexibility in defining how the collected monitoring data should be separated or aggregated, across a single appliance or multiple appliances, for display in the Tivoli Enterprise Portal. The following examples illustrate how data can be separated or aggregated for managing the data from various domains and appliances:

- **Separation of data at the domain:** You can view the services management data for the resources in a single domain, separate from the data for resources in other domains.
- **Aggregation of data across domains:** You can view the services management data for the resources in several domains (for example, all of the domains on a given DataPower appliance) in an aggregated form, with no regard for the domain in which individual resources are defined.
- **Separation of data at the appliance:** You can view the services management data for resources on a single DataPower appliance, separate from the data for resources on other appliances.

- **Aggregation of data across appliances:** You can view the services management data for the resources on several DataPower appliances (for example, all of the appliances in a load-balancing *cluster*) in an aggregated form, with no regard for what activity occurs on each individual appliance.

By default, the DataPower data collector aggregates data for all of the monitored domains on a single DataPower appliance (even if the domains are accessed using different credentials), and keeps the data from each DataPower appliance separated. See “Enabling the data collector” on page 21 for more information.

A single instance of the DataPower data collector can monitor any number of DataPower appliances, limited only by the memory, CPU power and other resources available to it. Note that IBM only supports running a single instance of the data collector on any given system.

Aggregation

For most IBM Tivoli Composite Application Manager for SOA data collectors, aggregation of data is performed for each application server runtime environment. The nodes in the Tivoli Enterprise Portal Navigator view represent individual application server runtime environments that have their own individual data collectors. Because the DataPower data collector can monitor multiple DataPower appliances and multiple domains within each appliance, this single application server runtime environment, single data collector model no longer applies.

Using the DataPower data collector, you assign names to groups of data, referred to as *display groups*. Each group of data is displayed as its own node in the Tivoli Enterprise Portal. Using these named display groups, you can configure the DataPower data collector to gather information from any domains on any DataPower appliances for aggregation and display.

By carefully managing the way in which you name these display groups, the DataPower data collector can separate or aggregate the data in different ways, such as isolating data specific to an individual domain on a single DataPower appliance, or aggregating data across several DataPower appliances into a single display group. This simple display group naming mechanism gives you great flexibility in the separation and/or aggregation of the data displayed in the Tivoli Enterprise Portal.

Deployment steps

To deploy the DataPower data collector in your environment, complete the following general steps as described in this document:

1. Install Sparkler 1 into your existing IBM Tivoli Composite Application Manager for SOA environment (see Chapter 2, “Installing Sparkler 1,” on page 5 for details).
2. Configure your DataPower appliances for monitoring (see “Configuring the DataPower appliance for monitoring” on page 20 for details).
3. Enable the DataPower data collector (see “Enabling the data collector” on page 21 for details).
4. Run the **startDPDC** script to start the data collector (see “Starting and stopping the data collector” on page 31 for details).

Configuring the DataPower appliance for monitoring

Before a DataPower appliance can be monitored by the DataPower data collector, configure the DataPower appliance as follows:

- Enable the XML Management Interface on the appliance.
- Check additional optional settings for each domain to be monitored.
- Configure a user account on the DataPower appliance for use with the DataPower data collector.

Enabling the XML Management Interface

The XML Management Interface on the appliance must be enabled using the DataPower administration console. To enable this service, complete the following steps:

1. Start the DataPower administration console in a Web browser (<https://hostname:9090/login.xml>).
2. Login to the administration console as **admin** in the default domain.
3. In the list at the left side of the console, navigate to **Objects → Network → XML Management Interface**.
4. On the **Main** tab, find the **WS-Management Endpoint** option and select the **on** radio button.
5. Click **Apply** to activate the changes and enable the WS-Management Endpoint.

Checking additional optional settings

For each domain to be monitored, check the following settings:

- In the list at the left side of the console, navigate to **Services → Miscellaneous → Web Services Agent**.
- Set **Admin State** to *enabled*.
- **Optional:** If you want the DataPower appliance to save metrics records when the DataPower data collector is not running, set **BufferMode** to *buffer*. With this setting, the appliance reports the saved metric data when the data collector starts.
- **Optional:** Adjust the values for **MaxRecords** and **MaxMemoryKB**
- **Optional:** If you want the data collector to record message content in addition to summary metrics, change **CaptureMode** from *faults* to *all-messages*.

Configuring a user account on the DataPower appliance

The DataPower user ID used by the data collector must belong to a user group with the following permissions:

- *Read* permission on the **Login XML-Mgmt** Resource Type in the default domain.
- *Read* permission on the **XML-mgmt** Resource Type in each domain to be monitored using this user ID.
- *Read* permission on the **(any)** Resource Type in each domain to be monitored using this user ID.

See your *DataPower WebGUI Guide* or *DataPower CLI Reference Guide* for details on configuring user group permissions.

Enabling the data collector

You should have already completed the installation procedure for installing Sparkler 1 in the existing IBM Tivoli Composite Application Manager for SOA environment on the computer where a supported application server runtime environment is installed. See “Installing Sparkler 1” on page 5 for details.

The DataPower configuration file

For the DataPower data collector, the KD4configDC command manipulates the contents of a special DataPower configuration file by adding sections to the file when new DataPower monitoring is enabled, and removing sections from the file when monitoring is disabled. Each invocation of the KD4configDC script is adding, updating or removing one section of the DataPower configuration file. Each section of the DataPower configuration file might be associated with its own data group, or it might be part of a larger data group to which other sections of the configuration file also belong.

The DataPower data collector uses this configuration file to identify the DataPower appliances that are to be monitored and to specify all of the information needed to communicate with those appliances. Typical information that is stored for each connection includes hostname and port, user ID and password, domains to monitor, and polling interval.

The configuration file is located in the `<TMS_DIR>/cma/kd4/config` directory and is called **KD4.dpdcConfig.properties**. This file is maintained separately from the existing configuration file, **KD4.dc.properties**. Here is a sample DataPower configuration file:

```
# Sample DataPower data collector configuration file
DataPower.count=3
#
DataPower.host.1=dpbox1
DataPower.port.1=5550
DataPower.path.1=/
DataPower.poll.1=60
DataPower.user.1=admin
DataPower.encpswd.1=#$*%&
DataPower.domainlist.1=default,testdom1
DataPower.displaygroup.1=dpbox1
#
DataPower.host.2=dpbox2
DataPower.port.2=5550
DataPower.path.2=/
DataPower.poll.2=30
DataPower.user.2=user1
DataPower.encpswd.2=&*%$#
DataPower.domainlist.2=userdom1,userdom2,userdom3
DataPower.displaygroup.2=user_doms
#
DataPower.host.3=dpbox2
DataPower.port.3=5550
DataPower.path.3=/
DataPower.poll.3=30
DataPower.user.3=admin
DataPower.encpswd.3=*%$#&
DataPower.displaygroup.3=all_doms
```

In this example, there are three sections in the configuration file. The properties in each of the three sections provide all of the information needed to establish and manage a single connection or session with each DataPower appliance.

Change the information in this configuration file using the KD4configDC command. Using various combinations of parameters in the KD4configDC command, you can achieve different monitoring configurations to separate or aggregate data among domains and appliances. See “Considerations for configuration of DataPower monitoring” on page 26 for more information.

Before configuring the DataPower data collector using KD4configDC, consult with your local systems management planners to understand which domains on which DataPower appliances are to be monitored and how the data from these domains and appliances should be separated or aggregated for display in the Tivoli Enterprise Portal.

Running the KD4configDC command

The syntax of the KD4configDC script is similar to the syntax for other supported IBM Tivoli Composite Application Manager for SOA data collectors. To run the KD4configDC command, first navigate to the following location, depending on your operating system:

- For supported Windows operating systems, navigate to %CANDLE_HOME%\CMA\KD4\bin.
- For supported AIX operating systems, navigate to \$CANDLEHOME/aix433/d4/KD4/bin.
- For supported Solaris operating systems, navigate to \$CANDLEHOME/sol283/d4/KD4/bin.
- For supported Linux (on Intel) operating systems, navigate to \$CANDLEHOME/li6243/d4/KD4/bin.
- For supported Linux (on z/OS) operating systems, navigate to \$CANDLEHOME/ls3243/d4/KD4/bin.

Run the KD4configDC command as follows:

- For supported Windows operating systems:
KD4configDC [-enable | -disable] -env 8 <env specific parameters>
- For supported AIX, Solaris, and Linux operating systems:
./KD4configDC.sh [-enable | -disable] -env 8 <env specific parameters>

You can issue the following command to view the online help for KD4configDC:

- For supported Windows operating systems:
KD4configDC -h
- For supported AIX, Solaris, or Linux operating systems:
./KD4configDC.sh -h

Specifying KD4configDC parameters

The <env specific parameters> defined for the DataPower invocation of the KD4configDC command are a series of key and value pairs that define the necessary properties for the affected section of the DataPower configuration file. These key and value pairs, which you can specify in any order on the command line, are shown in Table 1 on page 23.

Table 1. DataPower key and value pairs for KD4configDC

Parameter	Optional / Required	Default value	Description
-host <hostname or IP address>	Required		Defines the DataPower appliance hostname or IP address. This hostname is used to establish a socket connection and is used as part of the Web address pointing to the DataPower appliance. This can be any length string, with no blank characters. See "Creating node names in Tivoli Enterprise Portal" on page 24 regarding possible truncation of this value in the node name.
-user <user ID>	Required		Defines the DataPower appliance authentication user. This user must be a valid user for the DataPower appliance defined by the -host parameter. See your DataPower documentation for information about creating and managing user IDs for DataPower appliances. See "Configuring a user account on the DataPower appliance" on page 20 for more information.
-pswd <password>	Optional	User is prompted, if necessary	Defines the DataPower appliance authentication password, entered in clear text (not encoded). This password must be valid for the user defined in the -user parameter, and must be valid for the DataPower appliance defined by the -host parameter. This password is automatically converted to an encoded (masked) form and is stored in the DataPower configuration file. See your DataPower documentation for information about creating and managing passwords for DataPower appliances.
-port <port number>	Optional	5550	Defines the DataPower appliance port number. This port number is used to establish a socket connection and is used as part of the Web address pointing to the DataPower appliance. This value must be an integer from 0 to 65535. If this parameter is not specified, the default value is used.
-path <path string>	Optional	/	Defines the DataPower appliance path. This path is used as part of the Web address pointing to the DataPower appliance. This can be any valid path string, with no blank characters.
-poll <polling interval>	Optional	60 seconds	Defines the DataPower appliance polling interval (in seconds). The data collector waits this amount of time between each poll of the DataPower appliance. This must be an integer value, specified in seconds, between 5 and 300 (5 seconds to 5 minutes).

Table 1. DataPower key and value pairs for KD4configDC (continued)

Parameter	Optional / Required	Default value	Description
<code>-domainlist</code> <i>domainA,</i> <i>domainB,</i> <i>...domainZ</i>	Optional	No domainlist property is generated	Defines the DataPower appliance domain list. This is a comma-separated list of domains to be monitored on the associated DataPower appliance. Any domains in this list that are not authorized to the user defined by the <code>-user</code> parameter are not monitored. Each domain can be any string, with no blank characters. If you specify more than one domain name, separated by commas, the entire domain list must be enclosed in double quotes (for example, <code>-domainlist "domain1, domain2, domain3"</code>). See "Considerations for configuration of DataPower monitoring" on page 26 for more information on using this domain list.
<code>-displaygroup</code> <i><display group></i>	Optional	No displaygroup property is generated	Defines the DataPower appliance display name. The name can be any string, with no blank characters, up to 16 characters long. See "Creating node names in Tivoli Enterprise Portal" regarding possible truncation of this value in the node name. See "Considerations for configuration of DataPower monitoring" on page 26 for more information on the use of this property.

Creating node names in Tivoli Enterprise Portal

The DataPower data collector host name is combined with the value of the `-displaygroup` parameter, separated by a hyphen, to form the node name seen in the Tivoli Enterprise Portal Navigator view.

For example, suppose that the KD4configDC is run with the `-displaygroup` parameter specified with a value of *dispName*:

```
KD4configDC -enable -env 8 -host appHost -user user1 -displaygroup dispName
```

If the DataPower data collector host name is *dpdcHost*, then the node name that is displayed in the Tivoli Enterprise Portal Navigator view is *D4:dpdcHost-dispName*.

If the `-displaygroup` parameter is not specified in the KD4configDC command, then the value of the `-host` parameter is used instead. The value of the `-host` parameter is combined with the DataPower data collector host name to form the node name that is displayed in the Tivoli Enterprise Portal Navigator view.

For example, suppose the KD4configDC is run without specifying the `-displaygroup` parameter:

```
KD4configDC -enable -env 8 -host appHost -user user1
```

In this case the node name that is displayed in the Tivoli Enterprise Portal Navigator view is *D4:dpdcHost-appHost*.

Truncation of the node name: If the length of the node name exceeds 20 characters (including the hyphen), the node name is truncated from the end, meaning that the `-displaygroup` or `-host` portion is truncated. You should be aware

of this limitation of the Tivoli Enterprise Portal and choose values for *-displaygroup* and *-host* of appropriate length, to avoid truncation and possible confusion between similar names.

Examples of enabling the data collector

This section provides several examples of the kinds of monitoring that you can configure by running the KD4configDC command using various combinations of parameters.

The simplest form of the command

When you run the KD4configDC script to enable monitoring of a DataPower appliance, the simplest format of the command is:

```
KD4configDC -enable -env 8 -host <host> -user <user>
```

For example, with a host name of *dpbox1* and a user name of *admin*, the command is specified as:

```
KD4configDC -enable -env 8 -host dpbox1 -user admin
```

The DataPower configuration file is searched for an existing section that matches the specified host and user, and that has no domain list and no display group properties defined.

- If a matching section is found, you are prompted for a password, and the matching section is updated only by encoding the password and storing it in the configuration file for that section.
- If a matching section is not found in the configuration file, you are prompted for a password. A new section of the configuration file is created, with the specified values for *-host*, *-user* and *-pswd* stored in the configuration file for that section.

Note that no values for *-domainlist* or *-displaygroup* parameters are defined, and other optional properties for the section are set to their default values.

Because no value for *-displaygroup* is specified, the host name *dpbox1* is used as part of the node name that is displayed in the Tivoli Enterprise Portal (for example, *D4:dpdchost-dpbox1*)

Specifying domain list or display group parameters

If the *-domainlist* or *-displaygroup* parameters are specified in the KD4configDC command, they become part of the matching criteria that determines whether the section already exists or must be created as new. If a new section must be created, the values provided for the *-domainlist* and *-displaygroup* parameters are stored in the new section of the configuration file.

For example, if you have two domains, *domainA* and *domainB* that you are monitoring on a DataPower appliance defined by a host named *host1* and authorized to user *userABC*, run the KD4configDC command as:

```
KD4configDC -enable -env 8 -host host1 -user userABC -domainlist "domainA,domainB"
```

Note that the list of domains is surrounded by quotation marks, and there are no blank characters in the domain list.

As another example, suppose you want to create a node name to be displayed in the Tivoli Enterprise Portal Navigator view that contains the display group name of *all_domains*. This node name would be in the form of *D4:dpdchost-all_domains*. Specifying these parameters on the KD4configDC command, with authorized user *userABC*, would be as follows:

```
KD4configDC -enable -env 8 -host host1 -user userABC -displaygroup all_domains
```

Changing existing values for the domain list or display group: To update an existing section of the configuration file by changing the value of either the `-domainlist` or `-displaygroup` parameter, you must first disable that section of the configuration file by running the `KD4configDC` command using the `-disable` parameter (see “Disabling the data collector” on page 30), specifying the current values for the `-domainlist` or `-displaygroup` parameters.

After the section is disabled, run the `KD4configDC` command again, using the `-enable` parameter and specifying the desired values for `-domainlist` and `-displaygroup`.

Specifying the password parameter

The values specified by the `-user` and `-pswd` parameters are used to perform an HTTP basic authentication over the HTTPS session. The user and password combination must be a valid user ID and password defined using the administration interface of the DataPower appliance. Refer to your DataPower documentation for more information about creating authorized user names and passwords.

If the `-pswd` password parameter is specified in the `KD4configDC` command, its value is immediately encoded and saved in the configuration file, and you are not prompted for the password. Any password that had previously been specified for an existing section (that matches the search according to the other parameters) is replaced by the new password value.

For example, if an existing section for host name `H100` specified an authorized user of `u2233`, there would already be an encoded value of the valid password stored in the configuration file. To change the password value to `x1y2z3`, issue the `KD4configDC` command as follows:

```
KD4configDC -enable -env 8 -host H100 -user u2233 -pswd x1y2z3
```

Specifying optional parameters

If any of the remaining optional parameters (`-port`, `-path`, or `-poll`) are specified as part of the `KD4configDC` command, their values are stored in the configuration file in either the existing or newly created section. These values are not used as part of the matching criteria for locating existing sections of the configuration file.

The data collector uses the values in the `-host`, `-port`, and `-path` properties to construct the URL in the form of `https://host:port/path` that is used as the target of the messages sent to the DataPower appliance.

Considerations for configuration of DataPower monitoring

As previously described, you can use variations of the parameters specified in the `KD4configDC` command to define properties that configure the monitoring of one or more DataPower appliances. The following sections illustrate several possible configurations.

Multiple connections to the same DataPower appliance

The DataPower configuration file supports multiple connections to the same DataPower appliance by defining multiple sections of the configuration file with the same host and port properties. There are several reasons why you might want to do this, but the most obvious reason is to allow the data collector to use more than one set of authentication credentials.

For example, a DataPower appliance might have three domains (*dom1*, *dom2* and *dom3*) and two users (*userA*, with access only to *dom1* and *dom2*, and *userB* with access only to *dom3*).

In this situation, neither *userA* nor *userB* can be used to access all of the domains on the appliance. To monitor all of the domains on the appliance, you must define two sections in the DataPower configuration file, one that specifies *userA* in one section and *userB* in the other section.

To define the two sections in the configuration file, run `KD4configDC` twice, once for each user, as follows:

```
KD4configDC -enable -env 8 -host dpHost1 -user userA -domainlist "dom1,dom2"  
KD4configDC -enable -env 8 -host dpHost1 -user userB -domainlist dom3
```

Restricting the list of domains being monitored

The `-domainlist` parameter can be used to restrict the list of domains being monitored for a given user. For example, if *userA* has access to *dom1* and *dom2*, but you are interested in monitoring only *dom1*, the `-domainlist` parameter can specify only *dom1* to instruct the data collector to not monitor domain *dom2*:

```
KD4configDC -enable -env 8 -host dpHost1 -user userA -domainlist dom1
```

For each section in the DataPower configuration file, the domains that are monitored depend on the domains to which the user in that section has authorization, and the list of domains specified in the `-domainlist` parameter for that section. If this property is not specified, then all domains that are authorized to the user for that section are monitored.

Separating and aggregating data with the `-displaygroup` parameter

The `-displaygroup` parameter gives you substantial control over how the information from DataPower appliances is separated or aggregated in the Tivoli Enterprise Portal views. Following are some examples of the ways that the `-displaygroup` parameter can be used to separate or aggregate data.

Separating data by host name or display group

Suppose that you have two DataPower appliances (with different users and domains defined on each) and you want to view the information from these two appliances separately under two different nodes in the Tivoli Enterprise Portal Navigation view.

Because there are two DataPower appliances, there must be at least two sections in the DataPower configuration file, one for each appliance hostname or IP address. Separation of the information from these two appliances can be done without using the `-displaygroup` parameter if the two sections of the configuration file have different values for the `-host` parameter. For example:

```
KD4configDC -enable -env 8 -host dpHost1 -user userA  
KD4configDC -enable -env 8 -host dpHost2 -user userB
```

Note that the node name that is displayed in the Tivoli Enterprise Portal is based on the value of the `-host` parameter if the `-displaygroup` parameter is not specified. As long as the two sections of the configuration file have different values for the `-host` parameter, IBM Tivoli Composite Application Manager for SOA separates the data between the two appliances and displays it under two different nodes (for example, *D4:dpdcHost-dpHost1* and *D4:dpdcHost-dpHost2*) in the Tivoli Enterprise Portal Navigator view.

Alternatively, you could use the `-displaygroup` parameter to specify a more meaningful string than by using the value of the `-host` parameter. As long as the values for `-displaygroup` for the two sections are different, the data is separated under different Tivoli Enterprise Portal Navigator nodes. For example:

```
KD4configDC -enable -env 8 -host dpHost1 -user userA -displaygroup host1Domains
KD4configDC -enable -env 8 -host dpHost2 -user userB -displaygroup host2Domains
```

Because the `-displaygroup` parameter is specified, the node name is created from those values and should be displayed in the Tivoli Enterprise Portal Navigator view as `D4:dpdcHost-host1Domains` and `D4:dpdcHost-host2Domains`. However, note also that in this case, the length of the node name exceeds 20 characters, so each node name is truncated, as `D4:dpdcHost-host1Dom` and `D4:dpdcHost-host2Dom`.

Separating data between domains by domain list

Suppose that you have a DataPower appliance with a single user, `userA`, and two domains, `dom1` and `dom2`.

You can configure the data collector to separate the information from these two domains under two different nodes in the Tivoli Enterprise Portal Navigation view by creating two sections in the DataPower configuration file. The first section specifies a `-domainlist` value of `dom1` and a `-displaygroup` value of, for example, `Dom1`, and the second section specifies a `-domainlist` value of `dom2` and a `-displaygroup` value of, for example, `Dom2`.

Example:

```
KD4configDC -enable -env 8 -host dpHost1 -user userA -domainlist dom1
  -displaygroup Dom1
KD4configDC -enable -env 8 -host dpHost1 -user userA -domainlist dom2
  -displaygroup Dom2
```

The different values for the `-displaygroup` value cause IBM Tivoli Composite Application Manager for SOA to create two display groups that map to two different Tivoli Enterprise Portal Navigation view nodes: `D4:dpdcHost-Dom1` and `D4:dpdcHost-Dom2`, and the `-domainlist` values restrict the data collected in each of those display groups to just the intended domain, even though `userA` happens to have access to both domains.

Aggregating data across multiple domains

Suppose that you have a DataPower appliance with a single user, `userA`, and two domains, `dom1` and `dom2`. If the `-displaygroup` value is not specified in the section of the configuration file, the node name is created from the `-host` parameter value, and the information for both `dom1` and `dom2` is aggregated under a single node (`D4:dpdcHost-dpHost`) in the Tivoli Enterprise Portal Navigation view.

Aggregating data across multiple domains under multiple nodes

To aggregate the data from domains `domA1` and `domB1` under one node name and the data from `domA2` and `domB2` under a different node name, configure four sections in the DataPower configuration file, as follows:

- user `userA`, domainlist `domA1`, displaygroup `Group1`
- user `userB`, domainlist `domB1`, displaygroup `Group1`
- user `userA`, domainlist `domA2`, displaygroup `Group2`
- user `userB`, domainlist `domB2`, displaygroup `Group2`

Example:

```

KD4configDC -enable -env 8 -host dpHost -user userA -domainlist "domA1"
  -displaygroup Group1
KD4configDC -enable -env 8 -host dpHost -user userB -domainlist "domB1"
  -displaygroup Group1
KD4configDC -enable -env 8 -host dpHost -user userA -domainlist "domA2"
  -displaygroup Group2
KD4configDC -enable -env 8 -host dpHost -user userA -domainlist "domB2"
  -displaygroup Group2

```

This causes IBM Tivoli Composite Application Manager for SOA to aggregate the information from *domA1* and *domB1* under the Tivoli Enterprise Portal Navigation node *D4:dpdcHost-Group1* and to aggregate the information from *domA2* and *domB2* under the Tivoli Enterprise Portal Navigation node *D4:dpdcHost-Group2*

Aggregating data across all domains on a single host with multiple users

Suppose that you have a DataPower appliance with two users (*userA* and *userB*) that each have access to different sets of domains (*userA* has access to domains named *domA1* and *domA2*, the other user, *userB*, has access to domains named *domB1* and *domB2*, respectively). To monitor all of these domains, the DataPower configuration file must have at least two sections, one with a *-user* parameter value of *userA* and one with a *-user* parameter value of *userB*, but each specifying the same value for the host name in the *-host* parameter.

Example:

```

KD4configDC -enable -env 8 -host dpHost -user userA
KD4configDC -enable -env 8 -host dpHost -user userB

```

If the *-displaygroup* parameter is not specified, the information from all four domains would be aggregated under the same Tivoli Enterprise Portal Navigation node, for example, *D4:dpdcHost-dpHost*.

Aggregating data for a cluster of appliances

You might have a collection of DataPower appliances that are identically configured because they are used in a load-balancing or fail-over situation. Because each appliance has its own hostname or IP address, there must be at least one section in the configuration file for each appliance to inform the data collector how to communicate with each appliance.

If the *-displaygroup* parameter is not used, then the information for each of these appliances is separated into different nodes in the Tivoli Enterprise Portal Navigator view (for example, *D4:dpdcHost-dpHost1*), because the *-host* parameter value is used by default.

To see how this *cluster* of appliances is performing as a whole, without regard to the performance of individual appliances, specify the same *-displaygroup* value (for example, *Clust1*) for every appliance. This causes the data collector to aggregate the data from all of these devices into a single display group, which is then displayed under a single Tivoli Enterprise Portal Navigation node (for example, *D4:dpdcHost-Clust1*)

Example:

```

KD4configDC -enable -env 8 -host dpHost1 -user userA -displaygroup Clust1
KD4configDC -enable -env 8 -host dpHost2 -user userA -displaygroup Clust1

```

Additional considerations

You can mix and match any of the various configurations described in this chapter, which gives you a great deal of flexibility in defining how to aggregate or separate data in the Tivoli Enterprise Portal views. This flexibility might result in some DataPower domains being specified in more than one section. This is not considered an error.

You might want one section defined to allow data for a domain to be displayed on its own under one Tivoli Enterprise Portal Navigator node, and another section (or sections) defined to show the data for that domain aggregated with other domains, on the same or on a different machine. This is possible, but be aware that this results in the data for duplicated domains being retrieved multiple times by the data collector (once for each section in which it is configured). In this situation performance might be affected.

By default, if no `-domainlist` parameter and no `-displaygroup` parameter are used at all in the DataPower configuration file, the DataPower data collector aggregates data for all of the monitored domains on a single DataPower appliance (even if the domains are accessed using different credentials), and keeps the data from each DataPower appliance separated.

Disabling the data collector

When you run the `KD4configDC` script to disable monitoring of a DataPower appliance, the simplest format of the command is:

```
KD4configDC -disable -env 8 -host <host> -user <user>
```

The DataPower configuration file is searched for an existing section that matches the specified host and user and that has no `-domainlist` parameter and no `-displaygroup` parameters defined. If a matching section is found, the section is removed from the file. If no matching section is found, an error message is issued and the file remains unchanged.

If the `-domainlist` or `-displaygroup` parameters are specified in the `KD4configDC` command, they become part of the matching criteria that determines whether the section exists and removed if found.

All other parameters (`-port`, `-path`, `-poll`, or `-password`) are not required and are ignored, if specified on the command line.

Communication between data collector and appliance

Approximately every ten minutes, the data collector queries the DataPower appliance to see if the domain list has changed. If the domain list has changed, the data collector updates its information as needed and begins monitoring new domains defined on the Data Power appliance.

Communication errors between the data collector and the DataPower appliance that prevent the collection of data are logged to the console (the command line session where the `startDPDC` script was executed) and to the operations log. Additional information concerning the communications between the data collector and DataPower appliance can be obtained by setting the operations logging level for the data collector from `error` to `info` using the `updateLogging` Take Action command from the Tivoli Enterprise Portal.

Optimizing performance

The data collector is configured by default to monitor all services and operations, but with no logging of message content. This configuration optimizes the performance of processing data by not having to handle message content when it is not needed. When you modify any monitor controls to specify a content logging setting other than *none*, all of the full message content for all services and operations is retrieved from the appliance, and the data collector applies the configured monitor control rules (*headers/body/full*) by service and by operation.

Limitations

You can control the monitoring to limit the collection of data to specific services and operations, but you cannot use filter controls to reject messages to specific services or operations. To filter (reject) messages, use the Policy Rules and Actions of the DataPower Web Services Proxy configuration.

The data collector does not generate meaningful correlators for messages flowing through the DataPower appliance, so this data will not render properly in the topology, sequence diagram and pattern views in the IBM Web Services Navigator.

The DataPower data collector creates metric, content, operation and trace log files, but because message filtering is not supported, it does not create any action log files.

Depending on your configuration, the data collector can create multiple metric and content log files to emulate the existence of several different display groups. Each display group is represented as a separate node in the Tivoli Enterprise Portal Navigator view. The log files for each display group hold the data for all of the DataPower appliances that are configured for that display group, which might span multiple DataPower domains and include data from multiple DataPower appliances.

Before attempting to delete an application domain from DataPower that is being monitored by the DataPower data collector, you must either run the `KD4configDC` command to disable monitoring of that application domain, or stop the data collector. Otherwise, when you attempt to delete the domain using the DataPower administration console, you receive an error and the domain cannot be deleted.

Starting and stopping the data collector

To start the DataPower data collector, open a command prompt window and run the **startDPDC** script. No support is provided to run this data collector as a Windows, AIX, Solaris, or Linux service.

Once started, the data collector monitors the console for commands entered by the user. The only supported commands are **stop**, **quit**, and **exit**, all of which initiate an orderly shutdown of the DataPower data collector. When one of these commands is entered, the data collector waits for all communication sessions to end before the process terminates.

While the data collector is running, you can use the `KD4configDC` script to update the DataPower data collector configuration file and change which DataPower appliances and domains are being monitored. You do not need to stop and restart the data collector to activate these changes. The running data collector detects the changed configuration file within 40 seconds and adjust its monitoring to reflect the updated configuration.

Troubleshooting

Communication failures

In the event of a communications failure, the data collector attempts to initialize itself again and reestablish communication with the target DataPower appliance after the next polling interval (the polling interval is specified by the `-poll` value. This process is repeated each polling interval until communication is successfully reestablished, or until the data collector is stopped. Setting the operations logging level to `info` gives you the best indication of status of the communications between the data collector and the appliance.

Synchronizing time between computer systems

For data to display properly in the Tivoli Enterprise Portal, the DataPower appliance, the data collector system and the Tivoli Enterprise Monitoring Server system must have their clocks synchronized within 5 minutes of each other, in terms of the UTC time they report.

Password problems

User IDs and passwords for DataPower appliances are created and maintained at the appliance. See the documentation that is provided with your DataPower appliance for information about creating and managing user IDs and passwords. When you have reset this information at the DataPower appliance, use the `KD4configDC` script to update the DataPower configuration file with the latest user ID and password values for the appropriate section of the file.

Logging and Tracing

The DataPower data collector uses the same operation and trace log files as other data collectors in the IBM Tivoli Composite Application Manager for SOA product, but rather than creating separate operation and trace files for each display group, it creates a single operation file and a single trace file for all display groups.

The view in the Tivoli Enterprise Portal continues to operate as though you can configure logging and tracing for each display group, but the settings for all DataPower display groups are considered when the settings for the single log or trace writer instance are initialized, as follows:

- If any display group has trace set to ON, the DataPower data collector records tracing information for all display groups to the trace log.
- The DataPower data collector uses the lowest severity log setting (*Info*, *Warning*, or *Error*) of all display groups. For example, if a display group has a log setting of *Info*, then all informational, warning, and error messages for all display groups are written to the operations log.

The DataPower data collector will not update the global configuration settings seen in the Tivoli Enterprise Portal to reflect the actual settings being used for the various display groups. As a result, the Tivoli Enterprise Portal might still show different log and trace settings for each display group, even though the data collector is using common settings for all display groups.

In addition to the operations log, the DataPower data collector will also issue messages to the console where the `startDPDC` script was run. All messages issued to the console are also recorded in the operations log.

Appendix A. Messages

KD4DP0001I The DataPower data collector is exiting.

Explanation: The data collector is exiting.

KD4DP0002E The data collector is unable to locate or parse a configuration file named *fileName* due to the following error: *errorMessage*.

Explanation: This error indicates a problem processing the DataPower data collector configuration file.

Administrator Response: Verify that the configuration file specified in the error message exists and has the correct format. Run the data collector configuration program again, if necessary.

KD4DP0003I The data collector is monitoring domain: *domainName* on DataPower appliance: *hostName*.

Explanation: The DataPower data collector has discovered a domain on a DataPower appliance and will start a management subscription for the domain.

System action: The data collector attempts to start a subscription for the domain and will begin retrieving Web services metric data from the domain.

Administrator Response: Modify the DataPower data collector configuration file to customize which DataPower domains are monitored.

KD4DP0004E The data collector could not retrieve a domain list for DataPower appliance: *hostName* using user ID: *userID*. Detailed error message: *errorMessage*

Explanation: The data collector was unable to list the domains on a DataPower appliance. Possible causes:

- The configured user ID was not defined on the DataPower appliance.
- The configured password was incorrect.
- A communications failure occurred while communicating with the DataPower appliance.

System action: The data collector will periodically retry listing the domains.

Administrator Response: Verify that the DataPower data collector configuration file contains the correct host names, user IDs and passwords for the devices to be monitored.

KD4DP0005I The data collector has stopped monitoring domain: *domainName* on

DataPower appliance: *domainName*.

Explanation: The DataPower data collector has cancelled the management subscription for the domain.

System action: The data collector cancels the subscription for the domain.

Administrator Response: Modify the DataPower data collector configuration file to customize which DataPower domains are monitored.

KD4DP0006W The data collector is running. Enter STOP to stop the data collector.

Explanation: The data collector is currently monitoring DataPower appliances. Enter STOP at the console to stop the data collector.

System action: The data collector continues to run until STOP is entered at the console.

Operator response: Enter STOP at the console to stop the DataPower data collector.

KD4DP0007W The *actionName* request to DataPower appliance: *hostName* using user ID: *userID* encountered an error and did not complete successfully. Detailed error or fault code: *faultCode*

Explanation: The data collector was unable to complete a request to a DataPower appliance and received an error indication in response to the request. Possible causes:

- The configured user ID was not defined on the DataPower appliance.
- The configured password was incorrect.
- The XML Management Interface is not configured on the DataPower appliance.
- A communications failure occurred while communicating with the DataPower appliance.

System action: The data collector will periodically retry the request.

Administrator Response: Verify that the DataPower data collector configuration file contains the correct host names, user IDs and passwords for the appliances to be monitored. Verify that the XML Management Interface has been configured correctly on the DataPower appliance.

KD4DP0008E A communications failure occurred while the data collector was sending a request to DataPower appliance: *hostName* using user ID: *userID*. Detailed error message: *errorMessage*

Explanation: The data collector was unable to

communicate with the DataPower appliance. Possible causes:

- The DataPower appliance is not in an operational state.
- An incorrect TCP/IP address, TCP/IP hostname or TCP/IP port is configured for this DataPower appliance.
- The DataPower XML Management Interface is not configured correctly.
- A communications configuration or other networking problem exists.

System action: The data collector will periodically retry communications with the DataPower appliance.

Administrator Response: Verify that the DataPower data collector configuration file contains the correct host names, user IDs and passwords for the devices to be monitored. Verify that the XML Management Interface has been configured correctly on the DataPower appliance.

KD4DP0009E The DataPower data collector was unable to initialize using configuration file *fileName*. The data collector is exiting. Detailed error message: *errorMessage*

Explanation: The data collector was unable to initialize successfully due to a configuration file error. Possible causes:

- The configuration file referenced in the error message does not exist.
- File permissions prevent the data collector from reading the configuration file.
- The format of the configuration file is not valid.

System action: The data collector displays an error message and exits.

Administrator Response: Verify that the DataPower data collector configuration file exists and is correctly formatted. Use the DataPower data collector configuration script to create a valid configuration file.

KD4DP0010E The value for property: *propName* in the data collector configuration file: *fileName* is not valid. The configuration for DataPower appliance with an index value of: *index* is ignored.

Explanation: The data collector encountered a value for a required property that is not valid. The configuration for the DataPower appliance using this property is ignored.

System action: The data collector continues processing DataPower configuration definitions.

Administrator Response: Verify that the DataPower data collector configuration file exists and contains all required values. Use the DataPower data collector

configuration script to create a valid configuration file.

KD4DP0011W The value for property: *propName* in the data collector configuration file: *fileName* is not valid. The data collector will use the default value for this property: *defaultValue*.

Explanation: The data collector encountered an incorrect value for a property that has a default value. The default value will be used instead.

System action: The data collector continues processing DataPower configuration definitions.

Administrator Response: Verify that the DataPower data collector configuration file exists and contains all required values. Use the DataPower data collector configuration script to create a valid configuration file.

KD4DP0012E The DataPower data collector was unable to write the configuration file *fileName* with new property values. Detailed error message: *errorMessage*.

Explanation: The data collector was unable to successfully update the configuration file. Possible causes:

- The configuration file referenced in the error message does not exist.
- File permissions prevent the data collector from writing the configuration file.
- The configuration file is in use by another process.
- An operating system I/O error has occurred.

System action: The data collector displays an error message and continues without updating the configuration file. The values in the configuration file are unchanged but might be different from those that the data collector uses to communicate with the DataPower system.

Administrator Response: Verify that the DataPower data collector configuration file exists and is correctly formatted. Verify that permissions on the configuration file allow the data collector process to write to the file.

KD4DP0013E Syntax: Java *className* propertiesFileName.

Explanation: An attempt was made to start the data collector, but the syntax was incorrect. The data collector will exit.

Operator response: Use the appropriate startup script to start the data collector.

KD4DP0014I Enter the password for user *userName* ; it will be readable as it is typed:

Explanation: The DataPower data collector configuration for this appliance requires that the password for the specified DataPower user be specified. The password is displayed in clear text as it is entered.

Operator response: Enter the password for the specified DataPower user to continue the configuration process.

KD4DP0015E No monitoring is enabled that matches the specified command parameters. The disable request is ignored.

Explanation: The DataPower data collector configuration command requested that monitoring for a DataPower appliance be disabled. No current sections of the DataPower data collector configuration file matched the attributes specified on the disable command. To match a section of the configuration file the host, user, displaygroup, and domainlist properties must match exactly.

System action: No DataPower appliances are removed from the configuration file.

Operator response: To disable the monitoring of a DataPower appliance, specify the host, user, displaygroup, and domainlist information which matches the appliance to be removed from the data collector configuration file.

KD4DP0016E Monitor enable failed due to the exception: *exceptionMsg*

Explanation: The DataPower configuration tool encountered a Java exception while attempting to enable monitoring of a DataPower appliance.

System action: No DataPower appliances are added to the configuration file.

Operator response:

- Try to determine the problem from the exception text and correct if possible.
- Retry the task in progress when the error occurred.
- If the error continues to occur, contact IBM Software Support.

KD4DP0017E Monitor disable failed due to the exception: *exceptionMsg*

Explanation: The DataPower configuration tool encountered a Java exception while attempting to disable monitoring of a DataPower appliance.

System action: No DataPower appliances are removed from the configuration file.

Operator response:

- Try to determine the problem from the exception text and correct if possible.
- Retry the task in progress when the error occurred.
- If the error continues to occur, contact IBM Software Support.

KD4DP0018E The DataPower data collector configuration file: *fileName* could not be read. Detailed error message: *errorMessage*.

Explanation: The configuration tool was unable to successfully update the configuration file. Possible causes:

- File permissions prevent the data collector from writing the configuration file.
- An operating system I/O error has occurred.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response: Verify that the DataPower data collector configuration file exists and is correctly formatted. Verify that permissions on the data collector configuration file allow the configuration tool process to write to the file. If the error continues to occur, contact IBM Software Support.

KD4DP0019E The DataPower data collector configuration file: *fileName* could not be written. Detailed error message: *errorMessage*.

Explanation: The configuration tool was unable to successfully read the configuration file. Possible causes:

- File permissions prevent the data collector from reading the configuration file.
- An operating system I/O error has occurred.
- Another process has the file locked.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response: Verify that the data collector configuration file exists and is correctly formatted. Verify that permissions on the configuration file allow the configuration tool process to read the file. If the error continues to occur, contact IBM Software Support.

KD4DP0020E The DataPower data collector configuration file: *fileName* contains a DataPower.count property that is not valid.

Explanation: The DataPower configuration tool was unable to convert the DataPower.count property to a numeric value.

System action: The configuration tool displays an

KD4DP0021E • KD4DP0028E

error message and exits without modifying the configuration file.

Operator response: Verify that the data collector configuration file contains a valid `DataPower.count` property. If the error continues to occur, contact IBM Software Support.

KD4DP0021E An incorrect number of command parameters was specified.

Explanation: The minimum number of parameters required to enable or disable a DataPower device were not specified on the command.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response: See the documentation for the `KD4configDC` command for the correct command syntax.

KD4DP0022E The command parameter: *invalidKeyword* is not a valid keyword for this command.

Explanation: The DataPower configuration tool encountered an unrecognized keyword.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response: See the documentation for the `KD4configDC` command for the correct command syntax.

KD4DP0023E The required parameter: *missingKeyword* was not specified.

Explanation: A required configuration parameter was not specified.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response: See the documentation for the `KD4configDC` command for the correct command syntax.

KD4DP0024E The keyword: *duplicateKeyword* was specified more than once.

Explanation: A duplicate configuration parameter was specified.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response: See the documentation for the `KD4configDC` command for the correct command syntax.

KD4DP0025E The value specified for keyword: *invalidValue* is not valid.

Explanation: The value specified for a configuration keyword was not valid.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response: See the documentation for the `KD4configDC` command for the correct command syntax.

KD4DP0026E The encoding of the specified password failed due to the exception *exceptionMsg*.

Explanation: The DataPower configuration tool encountered a Java exception while attempting to encode a password for a user ID.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response:

- Try to determine the problem from the exception text and correct if possible.
- Retry the task in progress when the error occurred.
- If the error continues to occur, contact IBM Software Support.

KD4DP0027E The password entered is not valid.

Explanation: The password entered is not valid. It was either zero bytes in length or contained embedded blanks.

System action: The configuration tool displays an error message and exits without modifying the configuration file.

Operator response: Retry the command with a valid DataPower user password.

KD4DP0028E The DataPower data collector encountered an initialization error and is exiting. Error details: *exceptionMsg*.

Explanation: The DataPower data collector was unable to initialize successfully and is exiting.

System action: The data collector issues an error message and exits.

Operator response:

- Try to determine the problem from the error text and correct if possible.
- Retry the task in progress when the error occurred.
- If the error continues to occur, contact IBM Software Support.

KD4DP0029E The DataPower data collector could not obtain an exclusive lock on the configuration lock file: *fileName*.

Explanation: The data collector or the data collector configuration tool was unable to successfully lock the configuration file in preparation to write to the file. Possible causes:

- File permissions prevent the data collector from locking the configuration lock file.
- An operating system I/O error has occurred.
- The file is in use by another operating system process.

System action: The configuration tool displays an error message and exits without modifying the configuration file. The data collector continues monitoring DataPower appliances without modifying the configuration file.

Operator response: Verify that no other processes have the lock file open for write access. If the error continues to occur, contact IBM Software Support.

KD4DP0030W The value for the required property: *propName* in the data collector configuration file: *fileName* is not valid or is missing. The data collector will attempt to reinitialize itself.

Explanation: The data collector encountered a missing or incorrect value for a property that does not have a default value. The data collector will attempt to reinitialize itself.

System action: The data collector waits for a short period of time and then reads the data collector configuration file again and attempt to reinitialize itself.

Administrator Response: Verify that the DataPower data collector configuration file exists and contains all required values. Use the DataPower data collector configuration script to create a valid configuration file.

KD4EX0000I An example message to set a precedent, optionally including one or more *variables* or *labels*.

Explanation: additional information about the cause of the message and the action the program has taken or will take.

System action: describes what the system has done or will do in response to the situation that caused the message.

Operator response: describes the corrective actions the operator should take.

Administrator Response: describes the corrective actions the administrator should take.

Programmer response: describes the corrective actions the programmer should take.

KD4WT0001I *class_name* active, build level *build_level*.

Explanation: The component identified in the message became active during initialization. That component was compiled at the specified build level. The build level data should only be required when requested by IBM Software Support.

Operator response: No user action required.

KD4WT0002E The following exception occurred:
Exception: *exception_code*
Stack trace: *stack_trace_data*

Explanation: See the exception code and trace data for details of the error.

Operator response:

- Try to determine the problem from the exception text and correct if possible.
- Retry the task in progress when the error occurred.
- If the error continues to occur, contact IBM Software Support.

KD4WT0003W The configuration file *filename* could not be loaded. The default configuration is used instead.

Explanation: See message.

System action: The system automatically uses its built-in default configuration.

KD4WT0004E The following error occurred while loading the default configuration:
errorMessage. The *componentName* is inactive.

Explanation: This error indicates a problem with the internal libraries for the data collector.

System action: No Web services transactions are monitored.

Administrator Response: Contact IBM Software Support to correct the problem with the product libraries.

KD4WT0005E Unable to write to the data file *file_name*.

Explanation: The data collector is unable to write to the specified file. For example, the disk might be full or the file permissions might prohibit the operation.

System action: The data collector is inactive until the problem is corrected. After the problem is corrected, the system automatically reactivates after a short delay.

Administrator Response: Examine the file system for problems with the specified file.

KD4WT0006E The data collector does not support the *server_name* server.

Explanation: This data collector relies upon a small set of vendor-specific features. The server you are using is not supported by this version of the data collector.

System action: The data collector becomes inactive.

Administrator Response: Use a supported server implementation. Refer to the product documentation for a list of supported environments.

KD4WT0007E The *componentName* component is unable to initialize successfully due to the following error: *errorMessage*

Explanation: See message.

System action: The component is inactive.

Administrator Response: Correct the error indicated in the message and restart the application server.

KD4WT0008E The output file *fileName* could not be renamed for rollover.

Explanation: When a certain file-size threshold is reached, the product rolls the specified file over to a backup copy and begins a new file. This error indicates that the basic file system operations involved in this task have failed.

System action: The product will continue accumulating data in the original file. After the problem is corrected, the product rolls the current file over to its backup and continues its normal work.

Administrator Response: Correct the problem with the files or file system as promptly as possible. Until this problem is corrected, the risk of filling the file system is greatly increased.

KD4WT0009E This WebSphere Application Server installation is missing a required APAR. Install PQ89492 or an equivalent APAR and restart your server.

Explanation: WebSphere APAR PQ89492 delivers function that is required for this product. You must install this APAR or an equivalent APAR for this product to operate properly.

System action: The data collector component ceases to collect information until the WebSphere runtime environment is upgraded with this APAR.

KD4WT0010E The data collector is unable to locate or parse a configuration property named *property_name*.

Explanation: The data collector uses a configuration file named <INSTALL>/KD4/config/KD4.dc.properties.

This file either does not contain the named property, or the value of that property does not conform to the required syntax. These properties are enumerated, starting from one, up to the number specified on the corresponding "count" property.

System action: The data collector ignores this property and continues to use the set of configuration properties that it is able to successfully locate and parse.

Administrator Response: Refresh the configuration display for the corresponding server and update the set of configuration parameters. If you have modified the configuration file manually, ensure the following:

- The "count" property indicates the correct value.
- The correct number of "control" properties are present.
- Each "control" property contains the correct number and sequence of semicolon-delimited tokens.
- The server names on these properties match the server configuration and that they are free from typographic errors.

KD4WT0011E The Web service message has been rejected due to filter control configuration. WSDL Port = *port*, WSDL Operation = *operation*, remote system = *remote_system*, rejected at = *local_system*.

Explanation: The AddFiltrCntrl and DelFiltrCntrl Take Actions are used to control this feature. Refer to the product documentation for details about this feature.

System action: Web services messages that match the defined filter control criteria continue to be rejected until this filter control criteria is removed using the DelFiltrCntrl Take Action.

Administrator Response: Review the filter control criteria defined in the Data Collector Filter Control Configuration view of the Service Management Agent node for the system where the message was rejected. Use the DelFiltrCntrl and AddFiltrCntrl Take Actions to modify these filter criteria as needed. See the documentation or online help for more information about these commands.

KD4WT0012I The IPv6 address = *address* was found in the intercepted message. If the corresponding Web services message was incorrectly passed, check the filter criteria for an equivalent address and replace it with this one.

Explanation: IPv6 addresses may take any of several forms of text representations. The data collector requires these addresses to match as text. Therefore two representations that refer to the same address, but in different formats, will not match, allowing messages

to pass the data collector's filter. This might or might not be the expected behavior.

Administrator Response: If you expected the Web services message to be rejected but it was allowed through, then add the address given in this message using the AddFltrCntrl Take Action.

Appendix B. Obtaining IBM software support for IBM Tivoli Composite Application Manager for SOA

If you have a problem with your IBM software, you want to resolve it quickly. This section describes the following options for obtaining support for IBM software products:

- “Searching knowledge bases”
- “Obtaining fixes”
- “Receiving weekly support updates” on page 42
- “Contacting IBM Software Support” on page 42

Searching knowledge bases

You can search the available knowledge bases to determine whether your problem was already encountered and is already documented.

Searching the information center

IBM provides extensive documentation that can be installed on your local computer or on an intranet server. You can use the search function of this information center to query conceptual information, instructions for completing tasks, and reference information.

Searching the Internet

If you cannot find an answer to your question in the information center, search the Internet for the latest, most complete information that might help you resolve your problem.

To search multiple Internet resources for your product, use the **Web search** topic in your information center. In the navigation frame, click **Troubleshooting and support ► Searching knowledge bases** and select **Web search**. From this topic, you can search a variety of resources, including the following:

- IBM technotes
- IBM downloads
- IBM Redbooks™
- IBM developerWorks®
- Forums and newsgroups
- Google

Obtaining fixes

A product fix might be available to resolve your problem. To determine what fixes are available for your IBM software product, follow these steps:

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/support>.
2. Click **Downloads and drivers** in the **Support topics** section.
3. Select the **Software** category.
4. Select a product in the **Sub-category** list.
5. In the **Find downloads and drivers by product** section, select one software category from the **Category** list.

6. Select one product from the **Sub-category** list.
7. Type more search terms in the **Search within results** if you want to refine your search.
8. Click **Search**.
9. From the list of downloads returned by your search, click the name of a fix to read the description of the fix and to optionally download the fix.

For more information about the types of fixes that are available, see the *IBM Software Support Handbook* at <http://techsupport.services.ibm.com/guides/handbook.html>.

Receiving weekly support updates

To receive weekly e-mail notifications about fixes and other software support news, follow these steps:

1. Go to the IBM Software Support Web site at <http://www.ibm.com/software/support>.
2. Click **My support** in the upper right corner of the page.
3. If you have already registered for **My support**, sign in and skip to the next step. If you have not registered, click **register now**. Complete the registration form using your e-mail address as your IBM ID and click **Submit**.
4. Click **Edit profile**.
5. In the **Products** list, select **Software**. A second list is displayed.
6. In the second list, select a product segment, for example, **Application servers**. A third list is displayed.
7. In the third list, select a product sub-segment, for example, **Distributed Application & Web Servers**. A list of applicable products is displayed.
8. Select the products for which you want to receive updates, for example, **IBM HTTP Server** and **WebSphere Application Server**.
9. Click **Add products**.
10. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
11. Select **Please send these documents by weekly email**.
12. Update your e-mail address as needed.
13. In the **Documents** list, select **Software**.
14. Select the types of documents that you want to receive information about.
15. Click **Update**.

If you experience problems with the **My support** feature, you can obtain help in one of the following ways:

Online

Send an e-mail message to erchelp@ca.ibm.com, describing your problem.

By phone

Call 1-800-IBM-4You (1-800-426-4968).

Contacting IBM Software Support

IBM Software Support provides assistance with product defects.

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as DB2® and WebSphere products that run on Windows, or UNIX® operating systems), enroll in Passport Advantage in one of the following ways:

Online

Go to the Passport Advantage Web site at http://www.lotus.com/services/passport.nsf/WebDocs/Passport_Advantage_Home and click **How to Enroll**.

By phone

For the phone number to call in your country, go to the IBM Software Support Web site at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region.

- For customers with Subscription and Support (S & S) contracts, go to the Software Service Request Web site at <https://techsupport.services.ibm.com/ssr/login>.
- For customers with IBMLink™, CATIA, Linux, S/390®, iSeries™, pSeries®, zSeries®, and other support agreements, go to the IBM Support Line Web site at <http://www.ibm.com/services/us/index.wss/so/its/a1000030/dt006>.
- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage Web site at <http://www.ibm.com/servers/eserver/techsupport.html>.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the *IBM Software Support Handbook on the Web* at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region for phone numbers of people who provide support for your location.

To contact IBM Software support, follow these steps:

1. “Determining the business impact”
2. “Describing problems and gathering information” on page 44
3. “Submitting problems” on page 44

Determining the business impact

When you report a problem to IBM, you are asked to supply a severity level. Therefore, you need to understand and assess the business impact of the problem that you are reporting. Use the following criteria:

Severity 1

The problem has a *critical* business impact. You are unable to use the program, resulting in a critical impact on operations. This condition requires an immediate solution.

Severity 2

The problem has a *significant* business impact. The program is usable, but it is severely limited.

Severity 3

The problem has *some* business impact. The program is usable, but less significant features (not critical to operations) are unavailable.

Severity 4

The problem has *minimal* business impact. The problem causes little impact on operations, or a reasonable circumvention to the problem was implemented.

Describing problems and gathering information

When describing a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently. To save time, know the answers to these questions:

- What software versions were you running when the problem occurred?
- Do you have logs, traces, and messages that are related to the problem symptoms? IBM Software Support is likely to ask for this information.
- Can you recreate the problem? If so, what steps were performed to recreate the problem?
- Did you make any changes to the system? For example, did you make changes to the hardware, operating system, networking software, and so on.
- Are you currently using a workaround for the problem? If so, be prepared to explain the workaround when you report the problem.

Submitting problems

You can submit your problem to IBM Software Support in one of two ways:

Online

Click **Submit and track problems** on the IBM Software Support site at <http://www.ibm.com/software/support/probsub.html>. Type your information into the appropriate problem submission form.

By phone

For the phone number to call in your country, go to the contacts page of the *IBM Software Support Handbook* at <http://techsupport.services.ibm.com/guides/contacts.html> and click the name of your geographic region.

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround that you can implement until the APAR is resolved and a fix is delivered. IBM publishes resolved APARs on the Software Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Appendix C. Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other 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.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation
Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106, Japan

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 might 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 the 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 (l) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX 78758 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 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.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples 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.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs.

Trademarks

AIX, Candle[®], DB2, developerWorks, eServer, IBM, the IBM logo, IBMLink, iSeries, Lotus, OMEGAMON, Passport Advantage, pSeries, Rational, Redbooks, S/390, Tivoli, Tivoli Enterprise[™], the Tivoli logo, WebSphere, z/OS, and zSeries are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both.

Microsoft[®], Windows, and the Windows logo are registered trademarks of Microsoft Corporation in the United States, other countries, or both.

Intel is a registered trademark of Intel Corporation or its subsidiaries in the United States and other countries.

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

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.



Java and all Java-based trademarks and logos are trademarks or registered 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.

Index

A

- agent 1
- aggregation 19
 - across multiple domains 28
 - all domains, single host, multiple users 29
 - cluster of appliances 29
 - DataPower 17, 27
 - planning for 18
 - using display group 27
- appliance
 - aggregating multiple, DataPower 17
 - cluster, DataPower 29
 - communicating with data collector 30
 - DataPower 17
- application domain
 - DataPower 18
- application server runtime environment 1
- asynchronous flows
 - in SCA 13

B

- BEA WebLogic Server 2
- books
 - see publications 3

C

- cluster
 - DataPower appliances 29
- configuration file
 - DataPower 17, 21
- customer support
 - See Software Support

D

- data collector 1
 - communicating with appliance 30
 - DataPower 17
 - DataPower, deploying 19
 - DataPower, starting 31
 - DataPower, stopping 31
 - disabling for DataPower 30
 - disabling for JBoss 11
 - disabling for SCA 16
 - enabling for DataPower 21
 - enabling for JBoss 9
 - enabling for SCA 13
 - IBM WebSphere Application Server, in WESB and WPS 15
 - JBoss 9
 - multiple enabled 5
 - multiple, in same environment 15
 - SCA 13
 - support, installing 7
- Data Collector Adapter 1

- DataPower 2, 17
 - appliance, cluster 29
 - appliance, configuring 20
 - appliance, polling 17
 - configuration file 21
 - configuring optional settings 20
 - data collector, starting and stopping 31
 - disabling data collector 30
 - enabling data collector 21
 - firmware upgrade 17
 - multiple connections 26
 - proxy 17
 - user account, configuring 20
 - user password 26
- developerWorks, searching for problem resolution in IBM 41
- display group
 - aggregation and separation 27
 - changing current value 26
 - DataPower 19
- displaygroup
 - in node name, DataPower 24
- domain
 - aggregating multiple, DataPower 17
 - application, DataPower 18
- domain list
 - changing current value 26
 - restricting user access 27

F

- filtering
 - SCA 15
- firmware
 - DataPower requirement 17
- fixes, obtaining 41

G

- getting help 4

H

- host
 - in node name, Datapower 24

I

- IBM developerWorks, searching for problem resolution in 41
- IBM Passport Advantage 2
- IBM Redbooks, searching for problem resolution in 41
- IBM software support 4
- IBM technotes, searching for problem resolution in 41
- IBM Web Services Navigator 2
- IBM WebSphere Application Server 2
 - and SCA data collector 15

- IBM WebSphere Application Server *(continued)*
 - communicating between SCA modules 15
- IBM WebSphere Enterprise Service Bus 1, 13
- IBM WebSphere Integration Developer 13
- IBM WebSphere Process Server 13
- information centers, searching for problem resolution 41
- installation 5
 - data collector support 7
 - distributed, order of 5
 - Tivoli Enterprise Monitoring Server 6
 - Tivoli Enterprise Portal Server 6
- Intelligent Remote Agent 1
- Internet
 - searching for problem resolution 41

J

- JAX-RPC 2
 - JBoss 9
- JBoss 2, 9
 - disabling data collector 11
 - enabling data collector 9

K

- KD4configDC
 - DataPower, optional parameters 26
 - for DataPower 22
 - for JBoss 9
 - for SCA 13
 - parameters, DataPower 22
- knowledge bases, searching for problem resolution 41

L

- limitations
 - DataPower 31
- Linux support 2, 5, 6
- logging
 - DataPower 32
 - SCA 15

M

- manuals
 - see publications 3
- mediation 13
- message
 - filtering 2
 - logging 2
- Microsoft .NET 2
- monitor control
 - SCA 15
- monitoring agent 1
- multiple connections
 - DataPower 26

N

- node name
 - in TEP, for DataPower 24

O

- OMEGAMON 2

P

- parameters
 - KD4configDC for DataPower, optional 26
- password
 - DataPower, specifying 26
- performance
 - DataPower 31
- poll
 - DataPower appliance 17
- problem determination
 - describing problems 44
 - determining business impact 43
 - submitting problems 44
- proxy
 - DataPower 17
- publications
 - feedback 3
 - online 3
 - ordering 3

R

- Redbooks, searching for problem resolution in IBM 41

S

- SCA 2, 13
 - communicating between SCA modules 15
 - components 13
 - disabling data collector 16
 - enabling data collector 13
 - filtering 15
 - IBM Web Services Navigator 16
 - limitations 16
 - logging 15
 - modules 13
 - monitor control 15
 - with IBM WebSphere Application Server data collectors 15
- separation
 - between domains by domain list 28
 - by display group 27
 - by host name 27
 - DataPower 27
 - using display group 27
- Service Component Architecture 13
- service oriented architecture 1
- services component architecture 2
- SOA 1
- Software Support
 - contacting 42

Software Support (*continued*)
 describing problems 44
 determining business impact 43
 receiving weekly updates 42
 submitting problems 44
sparkler 2
start DPDC
 starting DataPower data collector 19
startDPDC 30, 31
support information 4
synchronous flows
 in SCA 13

T

technotes, searching for problem resolution in IBM 41
Tivoli Enterprise Monitoring Server 6
Tivoli Enterprise Portal
 display group, DataPower 19
 node names, DataPower 24
Tivoli Enterprise Portal Server 6
tracing
 DataPower 32
transaction log service
 DataPower 17
troubleshooting 32
 communication 32
troubleshooting
 appliance password 32
 synchronizing time 32

U

updateLogging Take Action 30

W

WESB 1

X

XML Management Interface
 enabling, for DataPower 20



Printed in USA