IBM WebSphere InterChange Server

System Administration Guide

Version 4.3.0

IBM WebSphere InterChange Server

System Administration Guide

Version 4.3.0

Note!

Before using this information and the product it supports, read the information in “Notices” on page 163.

30September2004

This edition of this document applies to IBM WebSphere InterChange Server (5724-178), version 4.3.0 and IBM WebSphere Business Integration Toolset (5724-177), version 4.3.0, and to all subsequent releases and modifications until otherwise indicated in new editions.

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

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

Contents

About this document v
Audience v
Related documents v
Typographic conventions vi

New in this release vii
New in release 4.3 vii
New in release 4.2.2 vii
New in release 4.2.1 viii
New in release 4.2.0 viii
New in release 4.1.1 x
New in release 4.1.0 x
New in release 4.0.1 x
New in release 4.0.0 x

Chapter 1. Monitoring the system. 1
Using System Monitor 1
Using System Manager to monitor the system 24
Using the SNMP agent 28
Using WebSphere Business Integration Monitor 40

Chapter 2. Administering components of the system 47
Overview of administering the system 47
Administering InterChange Server. 48
Using System Manager 53
Administering connectors 57
Administering JMS transport optimization 70
Administering collaboration objects 72
Administering maps 78
Administering relationships 81
Using Relationship Manager. 85
Scheduling jobs 99
Backing up system components 104
Using repos_copy 108
Administering end-to-end privacy 119
Administering role-based access control (RBAC). 123

Chapter 3. Administering problem scenarios 131
Administering failed events 131
Administering run-time properties 147
Administering High-Availability (HA) systems 148
Administering the Object Request Broker 150

Appendix A. WebSphere MQ reference 155

Appendix B. Requirements for restarting IBM WebSphere Business Integration system components 157
InterChange Server restart requirements 157
Collaboration template restart requirements 157
Collaboration object restart requirements 158
Connector restart requirements 158
Business object restart requirements 159
Map restart requirements 160
Relationship restart requirements. 160

Notices . . . 163
Programming interface information . . . 164
Trademarks and service marks . . . 165

Index . . . 167

About this document

IBM(R) WebSphere(R) InterChange Server Version 4.3 and its associated toolset are used with IBM WebSphere Business Integration Adapters to provide business process integration and connectivity among leading e-business technologies and enterprise applications.

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

Audience

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

Related documents

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

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

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

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

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

Note: Important information about the products documented in this guide might be available in Technical Support Technotes and Flashes issued after this document was published. These can be found on the WebSphere Business Integration Support Web site, http://www.ibm.com/software/integration/websphere/support/. Select the component area of interest and browse the Technotes and Flashes sections. Additional information might also be available in IBM Redbooks at http://www.redbooks.ibm.com/.

Typographic conventions

This document uses the following conventions:

`courier font`	Indicates a literal value, such as a command name, filename, information that you type, or information that the system prints on the screen.
bold	Indicates a new term the first time that it appears.
italic, *`italic`*	Indicates a variable name or a cross-reference.
blue outline	A blue outline, which is visible only when you view the manual online, indicates a cross-reference hyperlink. Click inside the outline to jump to the object of the reference.
{ }	In a syntax line, curly braces surround a set of options from which you must choose one and only one.
[ ]	In a syntax line, square brackets surround an optional parameter.
...	In a syntax line, ellipses indicate a repetition of the previous parameter. For example, `option[,...]` means that you can enter multiple, comma-separated options.
< >	In a naming convention, angle brackets surround individual elements of a name to distinguish them from each other, as in `<server_name><`*`connector_name`*`>tmp.log`.
/, \	In this document, backslashes (\) are used as the convention for directory paths. For UNIX installations, substitute slashes (/) for backslashes. All IBM WebSphere InterChange Server product pathnames are relative to the directory where the IBM WebSphere InterChange Server product is installed on your system.
UNIX/Windows	Paragraphs inside a box with a UNIX or Windows label indicate notes listing operating system differences.
`%text%` and *`$text`*	Text within percent (%) signs indicates the value of the Windows `text` system variable or user variable. The equivalent notation in a UNIX environment is *`$text`*, indicating the value of the *`text`* UNIX environment variable.
ProductDir	Represents the directory where the product is installed.

New in this release

New in release 4.3

30September2004

In the 4.3 release of the IBM WebSphere InterChange Server product, this guide has been updated with the following changes:

- Revised the section "Using the SNMP agent" on page 28 with new information.
- Added the new section "Administering end-to-end privacy" on page 119, including a high-level definition, explanation of using the new Security - Privacy tab and information on configuring keystores.
- Added the new section "Administering role-based access control (RBAC)" on page 123, including a high-level description, administration details, roles, importing and exporting information, and the new tabs in the System Manager.
- Added references to BiDi in support of repos-copy command.
- Added the following new command line references for repos_copy, as well as a description of how to use these command line options:
  - -xdi
  - -xdn
  - -nc
  - -xmsp
- Added the new section "Administering JMS transport optimization" on page 70, including an overview the JMS transport behavior, instructions on turning optimization on and off, and how to maximize the potential of optimization.

New in release 4.2.2

February 2004

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

- Revised the section "Using WebSphere Business Integration Monitor" on page 40 with additional information about configuring flow monitoring.
- Made minor revisions to the following sections:
  - "Steps for stopping the SNMP agent" on page 32
  - "Steps for setting automatic and remote restart for a connector" on page 67
  - "Administering High-Availability (HA) systems" on page 148
  - "Administering the Object Request Broker" on page 150

December 2003

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

- Replaced VisiBroker ORB with IBM ORB
- Revised the following sections:
  - "Administering failed events" on page 131 to reflect minor GUI (Graphical User Interface) changes

- "Steps for setting automatic and remote restart for a connector" on page 67 to reflect the change from VisiBroker ORB to IBM ORB

- Added the following sections:
  - "Using Relationship Manager" on page 85
  - Steps for configuring stack trace
  - "Using WebSphere Business Integration Monitor" on page 40
- Made the following System Monitor terminology changes:
  - Removed all instances of ″Windows-based System Monitor.″ The tool formerly referred to as the Windows-based System Monitor no longer exists as a separate tool. The functionality of the Windows-based System Monitor has been moved to the InterChange Server Component Management view of System Manager. See the section "Using System Manager to monitor the system" on page 24 for further details.
  - Removed all instances of ″Web-based System Monitor″ and replaced them with ″System Monitor.″ The tool formerly referred to as the Web-based System Monitor is now called System Monitor. There is only one tool named System Monitor.

New in release 4.2.1

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

- Revised the section "Using System Monitor" on page 1
- Added the section Using persistent monitoring.
- Added the following properties to the section Steps for configuring other InterChange Server logging and tracing parameters:
  - SLEEP_TIME
  - MAX_TRACE_WRITE_TRIES
- Added remote start functionality to the `OADAutoRestartAgent` standard connector property (see Table 14 on page 68)
- Made editorial improvements, including some reorganization of information

New in release 4.2.0

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

- "New features"
- "New organization of this guide" on page ix
- "New terminology" on page ix

New features

The following new features are covered in this guide:

- Web-based System Monitor: The Web-based System Monitor is a new, Web-based tool that allows you to monitor the IBM WebSphere InterChange Server system from the Web. It coexists with the Windows-based System Monitor, which is a revised version of the System Monitor provided in previous releases. For a full description of the differences between these two versions of System Monitor, and for instructions on using each tool, see "Using System Monitor" on page 1.
- System Manager: System Manager is the revised tool that was known as CrossWorlds System Manager (CSM) in previous releases (see "New terminology" on page ix). Descriptions of how to use System Manager to

perform certain tasks exist throughout this guide, particularly in “Using System Manager” on page 53, but for a thorough description of the revised functionality of System Manager, see the *System Implementation Guide*.

- Flow Manager: Flow Manager is a new tool that allows you to manage unresolved flows. For instructions on using Flow Manager, see “Administering failed events” on page 131.
- Flow control: Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. For instructions on using flow control, see the following sections:
  - “Steps for reviewing default monitors” on page 2
  - “Steps for creating additional monitors” on page 5
  - “Steps for configuring system-wide flow control” on page 55
  - “Steps for configuring flow control for connectors” on page 69
  - “Steps for configuring flow control for collaboration objects” on page 77
- Repos_copy: Many new arguments have been added to and removed from the repos_copy tool. For more information, refer to “Using repos_copy” on page 108.
- Optimized InterChange Server recovery
- Long-lived business processes: Refer to “Steps for reconfiguring the timeout attribute for long-lived business processing” on page 78.
- Dynamic updates: Refer to Appendix B, “Requirements for restarting IBM WebSphere Business Integration system components,” on page 157.

New organization of this guide

This guide has been reorganized in the following ways:

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

New terminology

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

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

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

New in release 4.1.1

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

- "Locale for repos_copy files" on page 119

New in release 4.1.0

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

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

New in release 4.0.1

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

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

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

New in release 4.0.0

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

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

Additional new features covered in other IBM CrossWorlds guides include:

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

New distribution of content of this guide

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

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

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

Chapter 1. Monitoring the system

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

"Using System Monitor"

"Using System Manager to monitor the system" on page 24

"Using the SNMP agent" on page 28

"Using WebSphere Business Integration Monitor" on page 40

Using System Monitor

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

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

"Setting up System Monitor"

"Steps for logging on to System Monitor" on page 12

"Overview of the System Monitor interface" on page 13

"Setting up views to monitor the system" on page 14

"Customizing data" on page 17

"Using persistent monitoring" on page 22

Setting up System Monitor

Before you begin using System Monitor, you must have the required web servers, client browsers, and other software installed on your system.

You also must decide whether you want to use the default monitors provided with System Monitor or if you want to create additional monitors using the Monitor Definition wizard. For example, you might want a monitor called System Overview, which displays status and start time of all system components. You create this monitor using the Monitor Definition Wizard, a tool opened from System Manager.

Requirements for System Monitor

System Monitor requires the following software, listed in Table 1:

Table 1. Required software for System Monitor

Supported Web servers	Software required on Web server	Supported browsers
A Web application server which supports JSP versions 1.1 or later, and servlets versions 2.2 or later, such as IBM WebSphere Application Server versions 5.0.2 with fixpack 4 or 5.1, or Tomcat versions 4.1.24 or 4.1.27 (using IBM JDK 1.4.2)	DB2 Client (if using DB2 for InterChange Server repository database)	• Microsoft Internet Explorer 5.5 SP2 or higher, with Adobe SVG Viewer 3.0 plug-in • Netscape 4.7x (only), with Adobe SVG Viewer 3.0 plug-in (on a Windows 2000 or Windows XP operating system only)

Steps for reviewing default monitors

Refer to Table 2 below for information you can use to determine if you want to use the default monitors included with the System Monitor.

Note: Table 3 on page 9 contains a description of the display options listed in Table 2.. The section entitled "Examples of display options" on page 9 contains samples of the display options listed in Table 2..

Table 2. Default monitors

Default monitor	Definition	Display options	Available operations when viewing monitor
System Overview	Overview of the current status of all major components of the system: collaborations, connectors, maps, and relationships	Table tree (a table with expandable nodes in the first column that display more rows)	• Start, stop, pause, and shut down a collaboration • Start, stop, pause, and shut down a connector • Restart a connector agent • Start and stop a map • Start and stop a relationship
Collaboration Statistics	Current status and statistics of all collaborations in the system: • Status • Start time • Total flows • Successful flows • Failed flows • Events in process • Queued events • Max concurrent events	Table	• Start, stop, pause, and shut down

Table 2. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor
Connector Statistics	Current status and statistics of all connectors: • Status • Start time • Total up time • Business objects received • Business objects sent • Agent status	Table	• Start, stop, pause, and shut down • Restart connector agent
Map Status	Status of all maps	Table	Start and stop
Relationship Status	Status of all relationships	Table	Start and stop
Server Statistics	Current statistics of the server: the number of failed and successful calls, events, and flows	Stacked bar	None
Database Connections	Current status of database connections: • Number of free connections • Number of active connections • Maximum number of connections • Peak number of connections	Table	None
Message Queues	Current status of message queues: • Current depth • Maximum depth configured	Table	None
Business Objects	Current statistics of the business objects for a particular connector: business objects sent and business objects received	Table	None
Connector Subscriptions	Current statistics of the subscriptions for a particular connector: • Collaboration object • Initiator	Table	None
Collaboration Events	Current statistics of collaboration events, which includes the following information: • Events in process • Queued events	Bar	None

Table 2. Default monitors (continued)

Default monitor	**Definition**	**Display options**	**Available operations when viewing monitor**
Historical Server Statistics	Server statistics for a specific period of time. Statistical information: • Successful calls • Failed calls • Total calls • Successful events • Failed events • Total events • Successful flows • Failed flows • Total flows Time intervals: • Start date • End date	Bar	None
Historical Server Flows	Flow statistics of the server for a specific period of time at certain time intervals. Statistical information: • Successful flows • Failed flows • Total flows Time intervals: • 15 min., 30 min., hourly, daily, weekly, or monthly • Start date • End date	Line	None
Historical Collaboration Flows Stack	Flow statistics of a particular collaboration for a specific period of time at certain time intervals. Statistics information: • Successful flows • Failed flows • Total flows Time intervals: • 15 min., 30 min., hourly, 4 hours, 12 hours, daily, weekly, or monthly • Start date • End date	Stacked bar	None

Table 2. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor
Historical Collaboration Flows Line	Flow statistics of a particular collaboration for a specific period of time at certain time intervals. Statistics information: • Successful flows • Failed flows • Total flows Time intervals: • 15 min., 30 min., hourly, 4 hours, 12 hours, daily, weekly, or monthly • Start date • End date	Line	None
Event Rate	Current number of processed events per minute	Meter	None
Flow Control	Current state of collaboration objects and connectors under Flow Control: • Buffered events • Max event capacity • Blocked status (does not apply to non-blocking collaboration) • Events pending in database (applies only to non-blocking collaborations) • Saturated status	Table	None
State Change Log	Current persisted state changes on a component for a specified time period. State change information: • Time stamp • State Time intervals: • Start date • End date	Table	None

Steps for creating additional monitors

Before you begin creating additional monitors, review the existing default monitors in Table 2 on page 2, to see if the monitor you want to create already exists.

Perform the following steps to create a monitor:

1. Open System Manager.

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

   Note: If the status does not change to **connected**, make sure the selected InterChange Server instance is running.
4. Right-click the server instance, then click **Monitor Definition Wizard**. The Monitor Definition Wizard appears. See Figure 1..

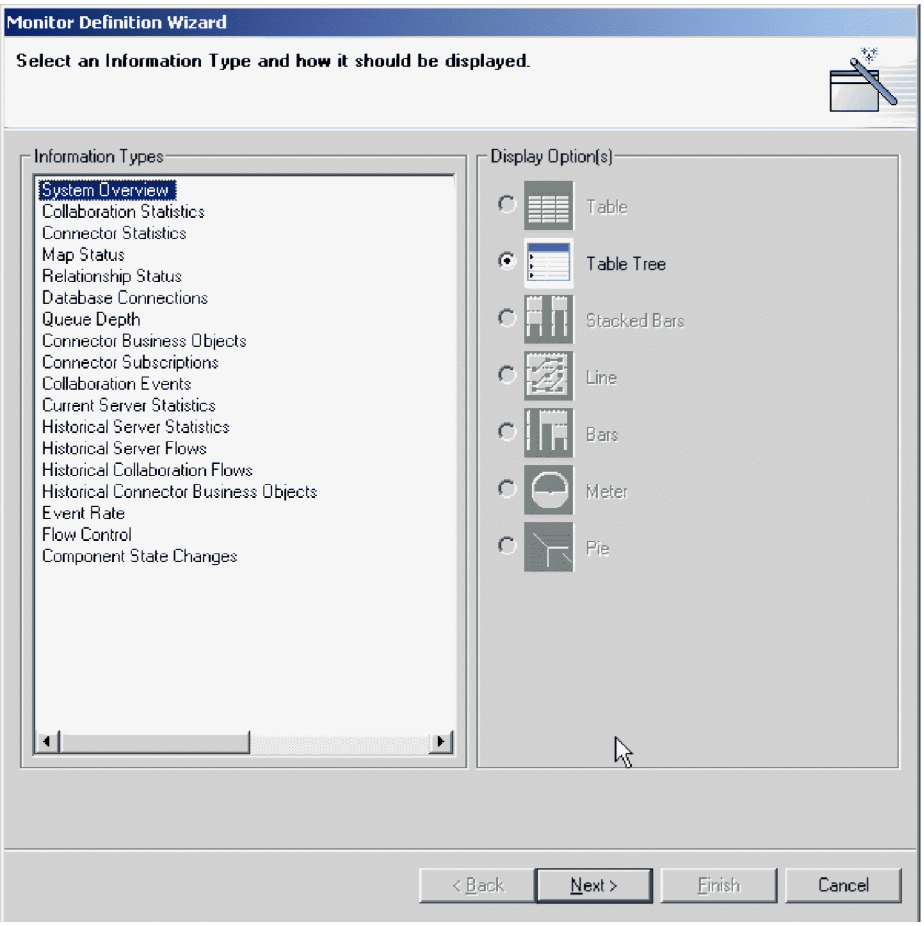

Monitor Definition Wizard

Select an Information Type and how it should be displayed.

Information Types

- System Overview
- Collaboration Statistics
- Connector Statistics
- Map Status
- Relationship Status
- Database Connections
- Queue Depth
- Connector Business Objects
- Connector Subscriptions
- Collaboration Events
- Current Server Statistics
- Historical Server Statistics
- Historical Server Flows
- Historical Collaboration Flows
- Historical Connector Business Objects
- Event Rate
- Flow Control
- Component State Changes

Display Option(s)

○ Table

● Table Tree

○ Stacked Bars

○ Line

○ Bars

○ Meter

○ Pie

Figure 1. Monitor Definition Wizard, page for selecting information type and display option

5. Select the type of information you want in the monitor from the **Information Types** list, and select how you want the information displayed under **Displayed Option(s)**.

   Each information type has one or more available display options, and each display option has configurable properties. When you select an information type, only the display options for that information type are available under **Displayed Option(s)**. For a description of the configurable properties of each display option, see “Steps for using monitor display options” on page 8, and for examples of how the display options appear in System Monitor, see “Examples of display options” on page 9.

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

6. Click **Next**. The Specify Monitor Properties page appears (see Figure 2).

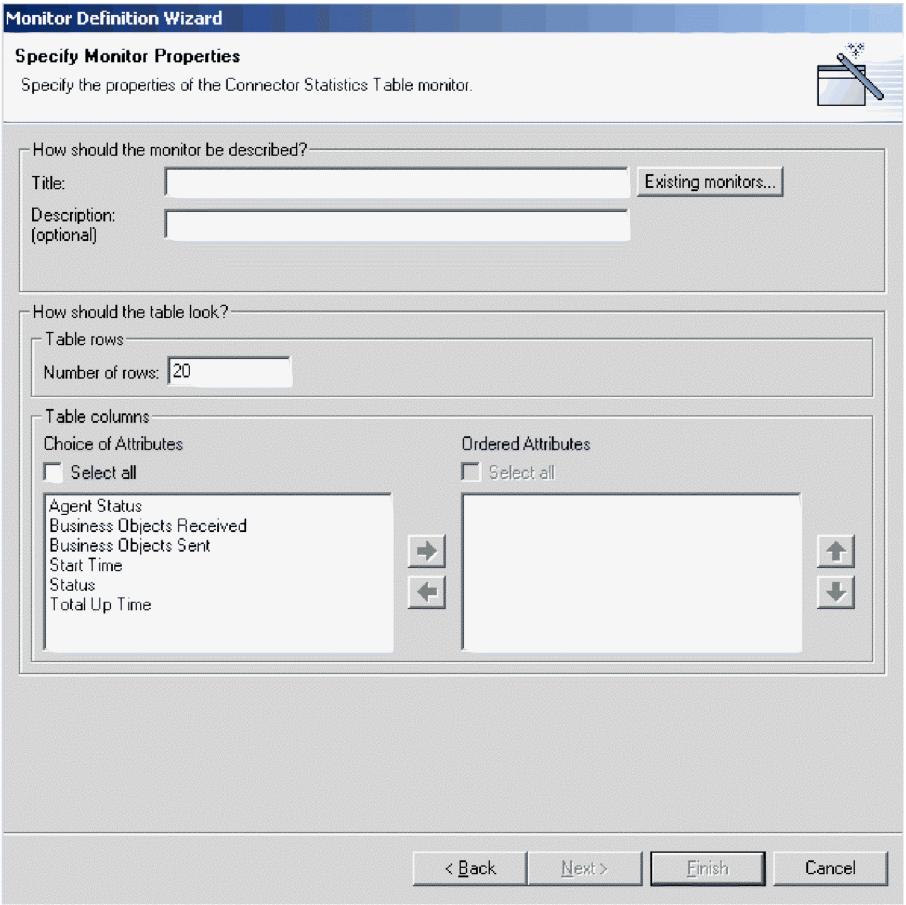

Monitor Definition Wizard

Specify Monitor Properties

Specify the properties of the Connector Statistics Table monitor.

How should the monitor be described?

Title:

Description: (optional)

How should the table look?

Table rows

Number of rows: 20

Table columns

Choice of Attributes

□ Select all

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

Ordered Attributes

□ Select all

Figure 2. Monitor Definition Wizard, Specify Monitor Properties page

7. Add the following information on the Specify Monitor Properties page:
   - Type a name for the new monitor in the **Title** field. To make sure you do not use an existing monitor name, click **Existing Monitors** to view a list of existing monitors.
   - (Optional) Type a description in the **Description** field.
   - Configure any additional properties available for the display option. These choices depend on the information type and display option that you chose on the previous page. For example, in Figure 2, you can type the number of rows to appear, select which attributes to include, and place the chosen attributes in a particular order. These options are available for both Table and Table Tree display options.

8. Do one of the following:
   - If, the attributes you chose can contain thresholds, the **Next** button is available. Click the **Next** button to configure the thresholds. The Specify Attribute Thresholds screen appears. For an example of a Specify Attribute Thresholds screen, see Figure 9 on page 24. In the Specify Attribute Thresholds screen, you can optionally type a numeric value in the threshold

field for each attribute. When running the monitor, if the value of an attribute exceeds the value of the threshold set for that attribute, the cell that contains the attribute value appears highlighted in the table.

- If the attributes don't contain thresholds, the **Finish** button is available. Click **Finish**. The following message appears: "The monitor was created successfully. Do you want to create another monitor?" Click **Yes** or **No**.

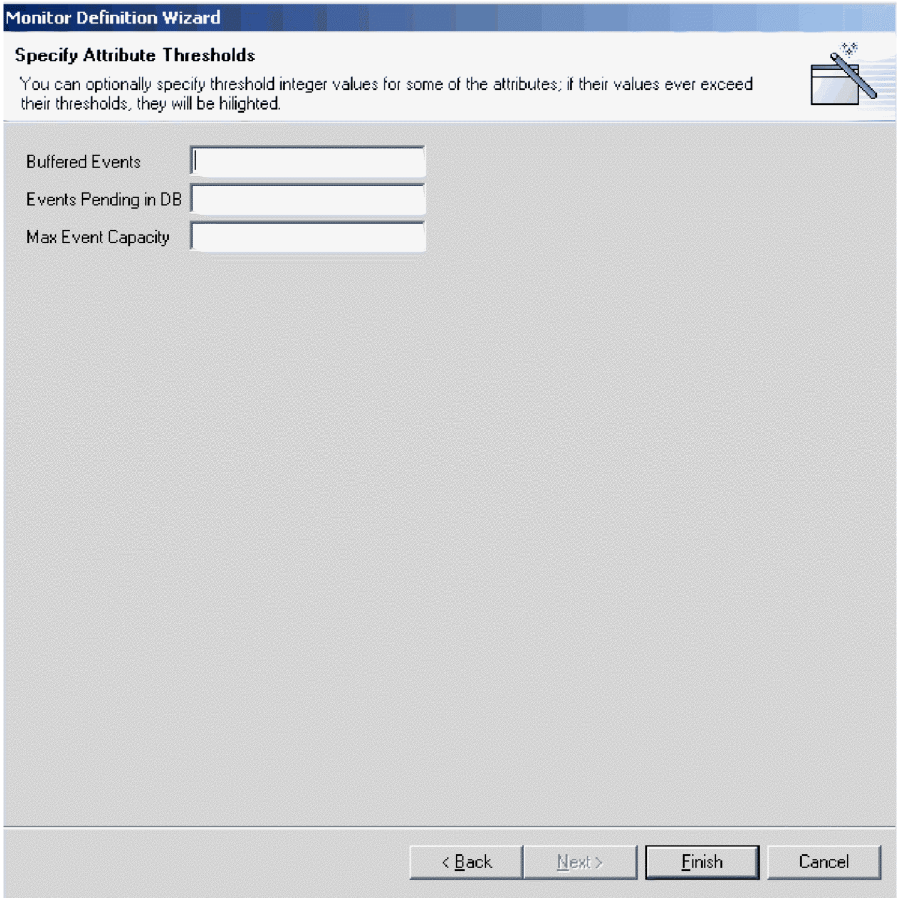

Monitor Definition Wizard

Specify Attribute Thresholds

You can optionally specify threshold integer values for some of the attributes; if their values ever exceed their thresholds, they will be hilighted.

Buffered Events

Events Pending in DB

Max Event Capacity

Figure 3. Monitor Definition Wizard, Specify Attribute Thresholds screen

Steps for using monitor display options

Perform this step to use display options for monitors you are creating with the Monitor Definition Wizard (see "Steps for creating additional monitors" on page 5) or using monitors in System Monitor (see "Steps for setting display properties for monitors" on page 18:

Refer to Table 3 to determine ways you can configure display options when creating monitors in Monitor Definition Wizard, or when you are using the monitors in System Monitor. (For examples of the display options, see "Examples of display options" on page 9.)

Table 3. Configurable display options for monitors

Display option	Properties you can configure when building monitors in Monitor Definition Wizard	Properties you can configure when using monitors in System Monitor
• Table • Table tree	• Columns to display • Order of columns • Number of rows to display	• Font and color settings of the labels and data • Number of rows to display
• Stacked bar • Line • Bar	None	• Font and color settings of the labels and data • Show or hide values
Meter	Meter threshold	Font and color settings of the labels and data

Examples of display options

The following exemplify the display options you can select when creating monitors in the Monitor Definition Wizard and how they appear in System Monitor:

- Table
- Table tree
- Line
- Bar
- Stacked bar
- Meter

Note: The data in the examples is not indicative of actual data in an InterChange Server system.

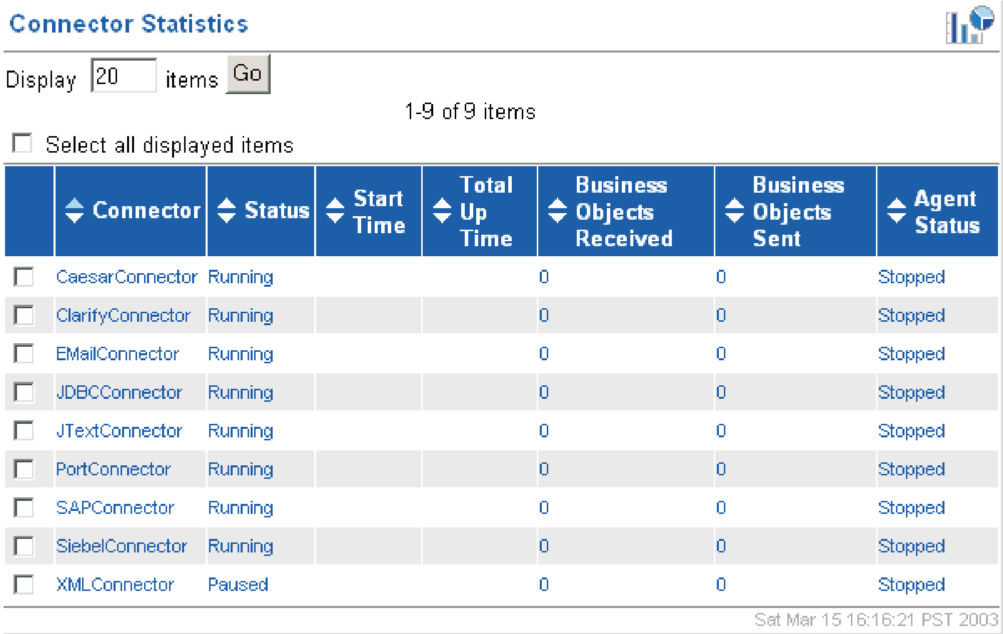

Connector Statistics

Display 20 items Go

1-9 of 9 items

□ Select all displayed items

	Connector	Status	Start Time	Total Up Time	Business Objects Received	Business Objects Sent	Agent Status
□	CaesarConnector	Running			0	0	Stopped
□	ClarifyConnector	Running			0	0	Stopped
□	EMailConnector	Running			0	0	Stopped
□	JDBCConnector	Running			0	0	Stopped
□	JTextConnector	Running			0	0	Stopped
□	PortConnector	Running			0	0	Stopped
□	SAPConnector	Running			0	0	Stopped
□	SiebelConnector	Running			0	0	Stopped
□	XMLConnector	Paused			0	0	Stopped

Sat Mar 15 16:16:21 PST 2003

Figure 4. Table display option

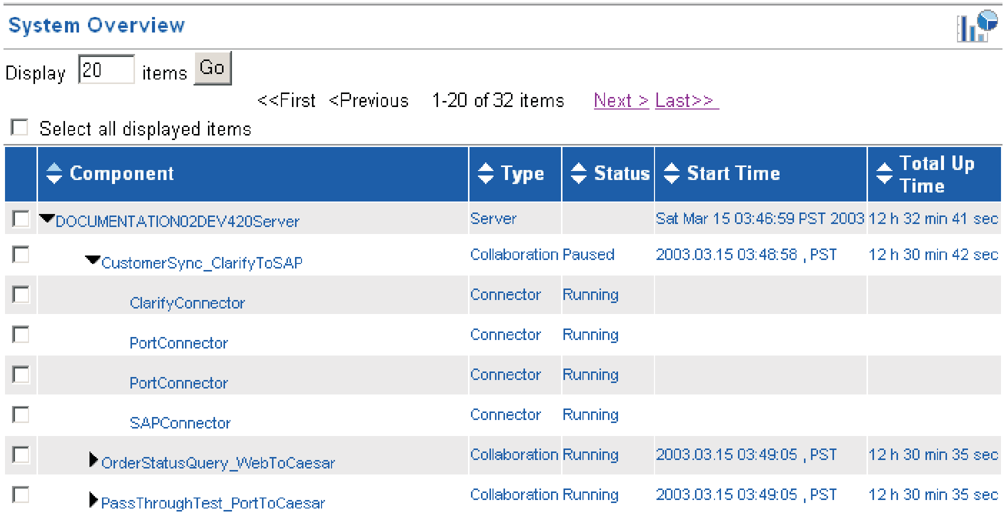

System Overview

Display 20 items Go

<<First <Previous 1-20 of 32 items Next > Last>>

□ Select all displayed items

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

Figure 5. Table tree display option

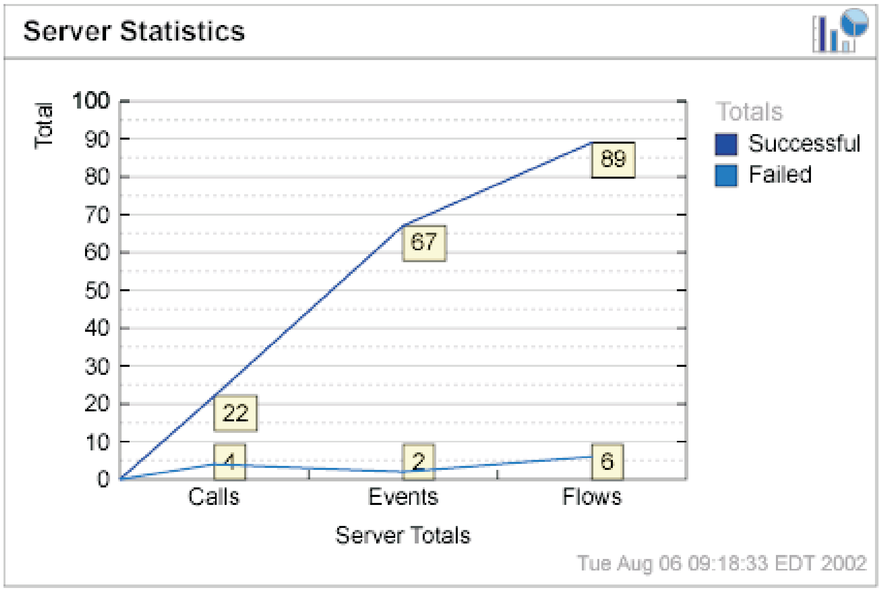

Figure 6. Line display option

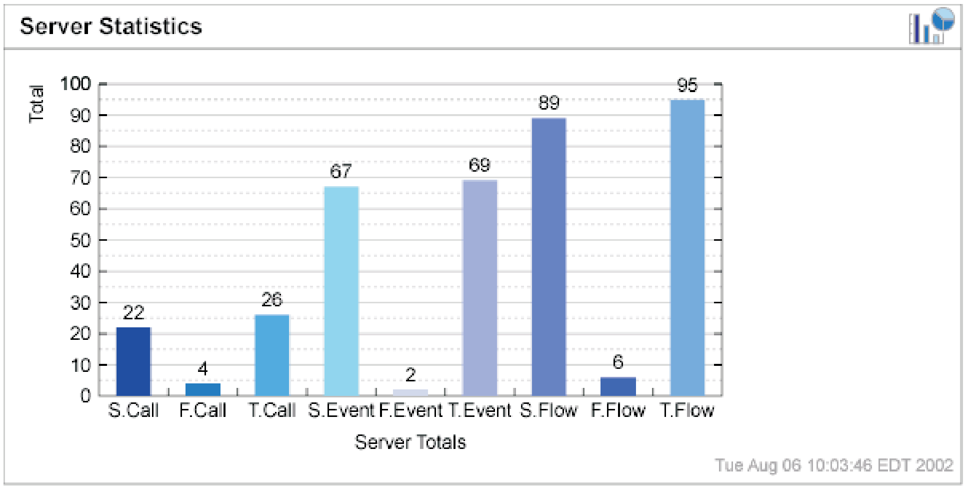

Figure 7. Bar display option

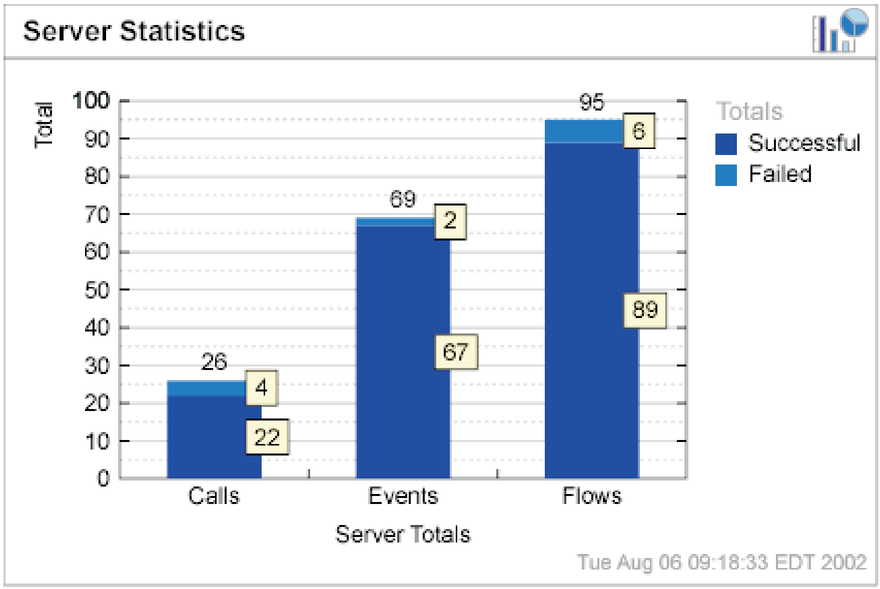

Figure 8. Stacked bar display option

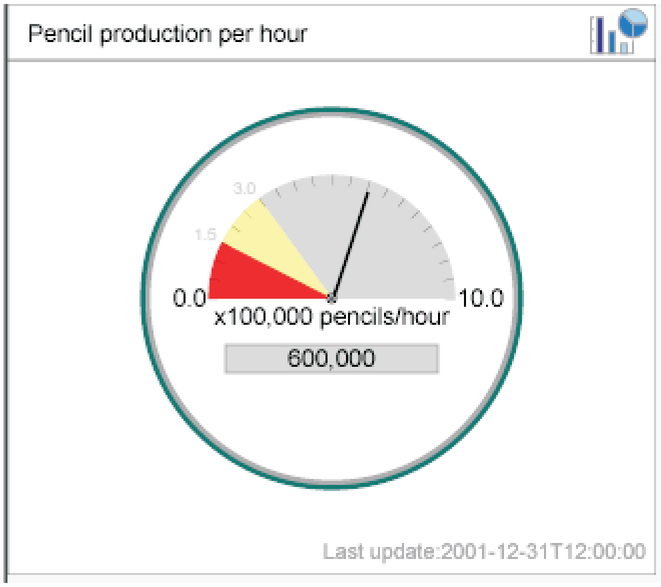

Figure 9. Meter display option

Steps for logging on to System Monitor

After you have either created new monitors or decided to use the default monitors, you are ready to log on to System Monitor to monitor the system.

Before you begin:

- Start InterChange Server on the machine being monitored.
- Make sure System Monitor and the application server are installed. For installation instructions, see *WebSphere InterChange Server Installation Guide*.
- Start the application server.
- Obtain the user name and password necessary for logging on to System Monitor. The user name and password are the same as those used when logging on to InterChange Server.

Perform the following steps to log on to System Monitor:

1. In a Web browser, navigate to the URL for System Monitor. The URL you use for System Monitor depends on whether you are using WebSphere Application Server or Tomcat. Refer to the *WebSphere Business InterChange Server Installation Guide* for additional information about setting up System Monitor to work with WebSphere Application Server or Tomcat.
   - If you are using WebSphere Application Server, the url is:
     ```
     http://HostName/ICSMonitor
     ```
     where *HostName* is the host name of the Web server machine.
   - If you are using Tomcat and did not change the port number, the URL is:
     ```
     http://HostName:8080/ICSMonitor
     ```
     where *HostName* is the host name of the Web server machine.

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

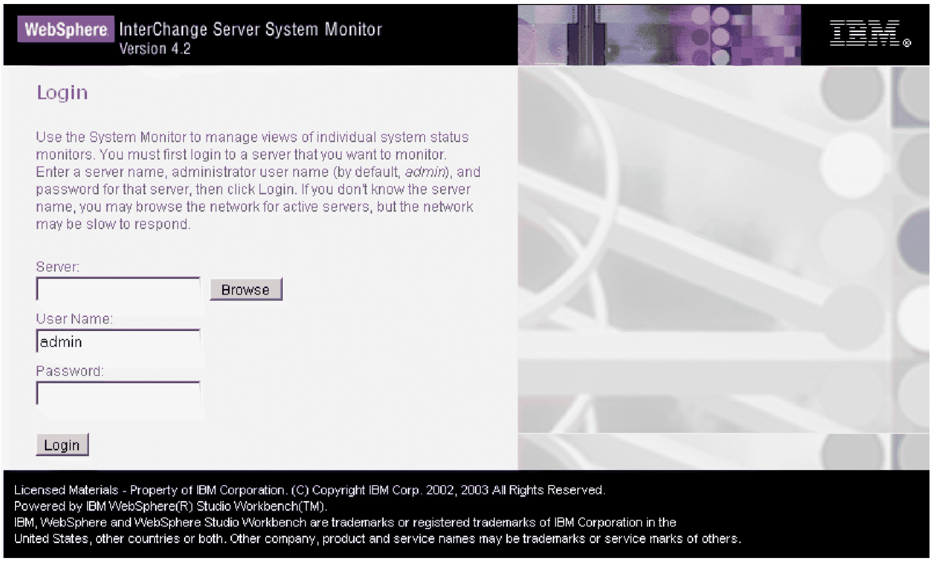

Login

Use the System Monitor to manage views of individual system status monitors. You must first login to a server that you want to monitor. Enter a server name, administrator user name (by default, *admin*), and password for that server, then click Login. If you don't know the server name, you may browse the network for active servers, but the network may be slow to respond.

Server:

User Name: admin

Password:

Licensed Materials - Property of IBM Corporation. (C) Copyright IBM Corp. 2002, 2003 All Rights Reserved.
Powered by IBM WebSphere(R) Studio Workbench(TM).
IBM, WebSphere and WebSphere Studio Workbench are trademarks or registered trademarks of IBM Corporation in the United States, other countries or both. Other company, product and service names may be trademarks or service marks of others.

Figure 10. System Monitor, Login window

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

   Note: If role-based access control is enabled, the user will not be allowed to log in to System Monitor unless they are assigned to a role that is granted permission to monitor the server. For more information on role-based access control, see "Administering role-based access control (RBAC)" on page 123.

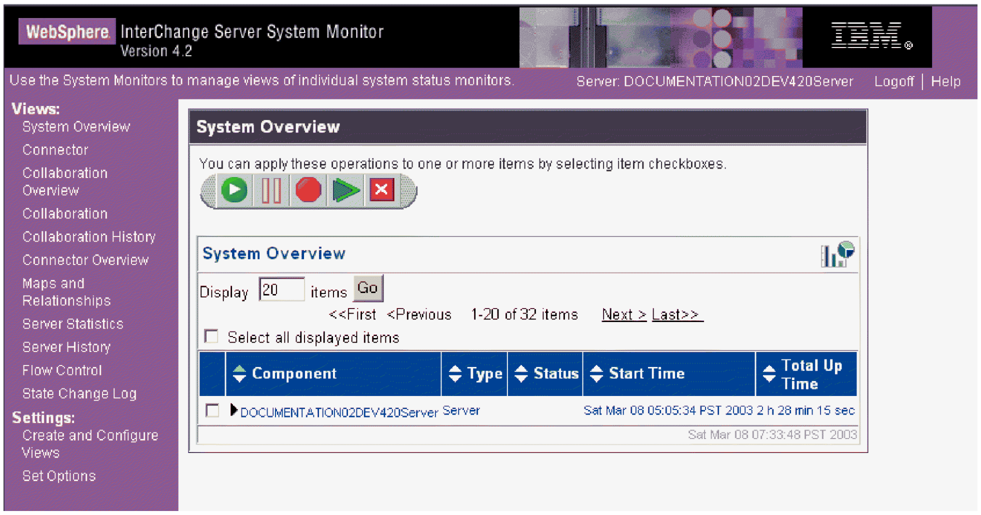

Use the System Monitors to manage views of individual system status monitors.

Server: DOCUMENTATION02DEV420Server

System Overview

You can apply these operations to one or more items by selecting item checkboxes.

System Overview

Display 20 items

1-20 of 32 items

□ Select all displayed items

	Component	Type	Status	Start Time	Total Up Time
□	DOCUMENTATION02DEV420Server	Server		Sat Mar 08 05:05:34 PST 2003	2 h 28 min 15 sec

Sat Mar 08 07:33:48 PST 2003

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

Overview of the System Monitor interface

System Monitor contains the following items:

- **List of views**: Initially, the views listed in the left column under **Views** are the default views provided with the installation of System Monitor, but you can add, change or delete views to suit your monitoring needs.
- **Create and Configure Views link**: This link opens the Create and Configure Views dialog box (see Figure 12 on page 16), which allows you to create, configure, or delete views. It also allows you to set the default view you see when you log on to System Monitor. For instructions on creating, configuring, or deleting views, or setting the default view you see when you log on, refer to the following sections:
  - "Steps for creating your own views" on page 16
  - "Steps for configuring views" on page 16
  - "Steps for deleting views" on page 17
  - "Steps for setting a default view" on page 18
- **Set Options link**: The **Set Options** link opens the Set Options dialog box (see Figure 14 on page 19), which allows you to do the following for system-wide or component settings:
  - Set the refresh rate of the views that display current statistics
  - Set the frequency of historical data captured for each component type
  - Reset component statistics to "0"
  - Capture component state changes
  - Delete component state change log
  - Delete historical statistics for all components in the system
  - Delete business object probe data log

  For instructions on using the Set Options dialog box, refer to the following sections:
  - "Steps for setting the refresh rate for run-time values" on page 19
  - "Steps for setting the frequency for historical data capture" on page 19
  - "Steps for resetting run-time statistic values" on page 20
  - "Steps for capturing state changes" on page 21
  - "Steps for deleting the state change log" on page 21
  - "Steps for deleting historical statistics" on page 21
  - "Steps for deleting the business object probe data file" on page 21
- **Default view**: A default view is displayed when you log on to System Monitor. The first time you open System Monitor, the System Overview view is displayed. To change the default view displayed, see "Steps for setting a default view" on page 18.
- **Logoff link**: The **Logoff** link allows you to log off System Manager.
- **Help link**: The **Help** link opens an HTML page with the following information:
  - A link to download the documentation set for the IBM WebSphere InterChange Server product
  - A directory location on your local machine where you can start the Table of Contents file with links to help topics. This assumes you have already downloaded the documentation set.

Setting up views to monitor the system

You can either begin monitoring the system using the default views, or you can add, change, or delete views before monitoring the system. The following sections describe how to use existing views or create and configure views from System Monitor. Views can contains one or more monitors. Several default views are included in the installation of System Monitor. You may use these default views or create new views. Before you can create and configure views, you must log on to System Monitor. For instructions on logging on to System Monitor, see “Steps for logging on to System Monitor” on page 12.

This section covers the following topics:

“Steps for using default views”

“Steps for creating your own views” on page 16

“Steps for configuring views” on page 16

“Steps for deleting views” on page 17

“Steps for resetting default views” on page 17

Steps for using default views

Perform the following step to use default views:

1. Open System Monitor.
2. In the left frame, select one of the views listed in Table 4 from the **Views** list.

   The table describes which monitor or monitors are contained in the view, and which display option is used. For descriptions of default monitors used in the views, see “Steps for reviewing default monitors” on page 2.

Table 4. Default views

Default view	**Monitor(s) and display options**
System Overview	System Overview monitor displayed in a table tree
Collaboration Overview	Collaboration Statistics monitor displayed in a table
Collaboration	• Collaboration Events monitor displayed in bar chart, and • Event Rate monitor displayed in a meter
Collaboration History	• Historical Collaboration Flows monitor displayed in a stacked bar chart • Historical Collaboration Flows monitor displayed in a line chart
Connector Overview	Connector Statistics monitor displayed in a table
Connector	• Business Objects monitor displayed in a table • Connector Subscriptions monitor displayed in a table
Maps and Relationships	• Map Status monitor displayed in a table • Relationship Status monitor displayed in a table
Server Statistics	• Server Statistics displayed in a stacked bar chart • Database Connections displayed in a table • Message Queues displayed in a table
Server History	• Historical Server Statistics displayed in a bar chart • Historical Server Flows displayed in a line chart
Flow Control	Flow Control monitor displayed in a table
State Change Log	State Change Log monitor displayed in a table

The table or chart for that view opens in the System Monitor main window. For examples, see "Examples of display options" on page 9.

Steps for creating your own views

Perform the following steps to create a view:

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

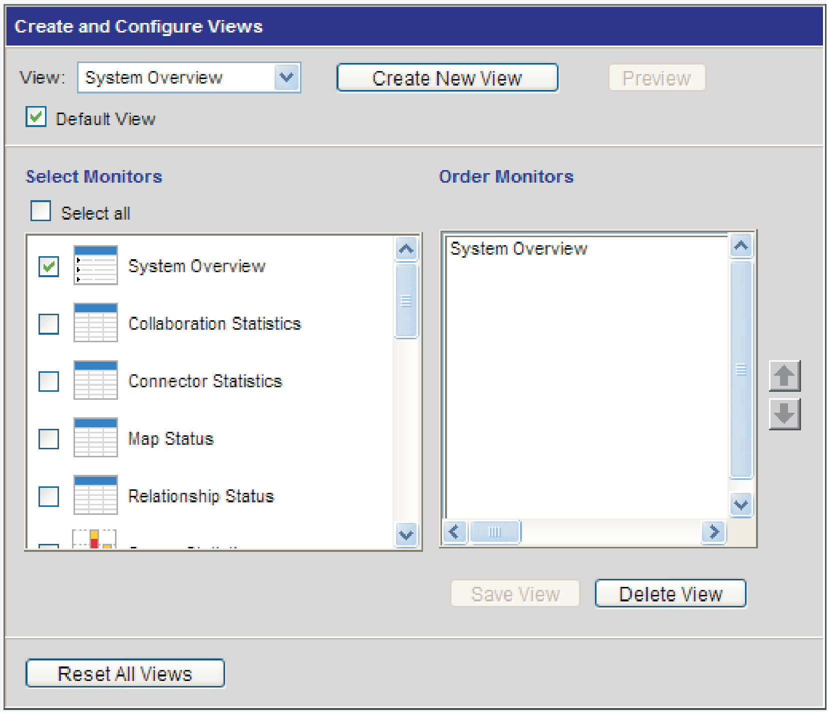

Figure 12. Create and Configure Views Window

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

Steps for configuring views

Perform the following steps to change an existing view:

1. Click **Create and Configure Views** in the left frame of System Monitor. The Create and Configure Views dialog box appears (see Figure 12).
2. Select the view you want to change from the **View** list.
3. Add monitors to or remove monitors from the view in the **Select Monitors** list. The revised monitors for the view appear in the **Order Monitors** list.

4. Use the up and down arrows to the right of the **Order Monitors** list to put the monitors in the order you want to view them.
5. Click **Preview** if you want to see a preview of the new view.
6. Click **Save View**. A “View was saved successfully” message appears.

Steps for deleting views

Perform the following steps to delete a view:

1. Click **Create and Configure Views** in the left frame of System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 16).
2. Select the view you want to delete from the **View** list.
3. Click **Delete View**. A message appears, asking if you are sure you want to delete the view.
4. Click **OK**. The view is removed from the Views list in the left frame of System Monitor.

Steps for resetting default views

After creating, deleting and configuring views, you may reset all views back to the original system defaults. All new or modified views will be lost when resetting back to the original defaults. Perform the following steps to reset default views:

1. Click **Create and Configure Views** in the left frame of System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 16).
2. Click the **Reset All Views** button. The Reset All Views pop-up window appears, displaying a listing of the number of views that will be deleted, reintroduced or modified.
3. Click **View Details** to view additional information on the changes that will occur if you proceed. When you have completed viewing details, click **Yes**. The views are set to the original system defaults. The navigation pane automatically updates the view listing.

Customizing data

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

“Steps for setting a default view” on page 18

“Steps for setting display properties for monitors” on page 18

“Steps for setting the refresh rate for run-time values” on page 19

“Steps for setting the frequency for historical data capture” on page 19

“Steps for resetting run-time statistic values” on page 20

“Steps for capturing state changes” on page 21

“Steps for deleting the state change log” on page 21

“Steps for deleting historical statistics” on page 21

“Steps for deleting the business object probe data file” on page 21

Steps for setting a default view

The default view is the view you first see when you log on to System Monitor.

Perform the following steps to change the default view:

1. Click **Create and Configure Views** from the left frame of System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 16).
2. Select the view that you want to be the default view from the **View** list.
3. Select the **Default View** check box.
4. Click **Save View**. A “View was saved successfully” message appears. The next time you log on to System Monitor, the view you selected as the default view is displayed.

Steps for setting display properties for monitors

The display options of monitors can be customized by changing the preferences of the display options.

Perform the following steps to change the appearance of a monitor:

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

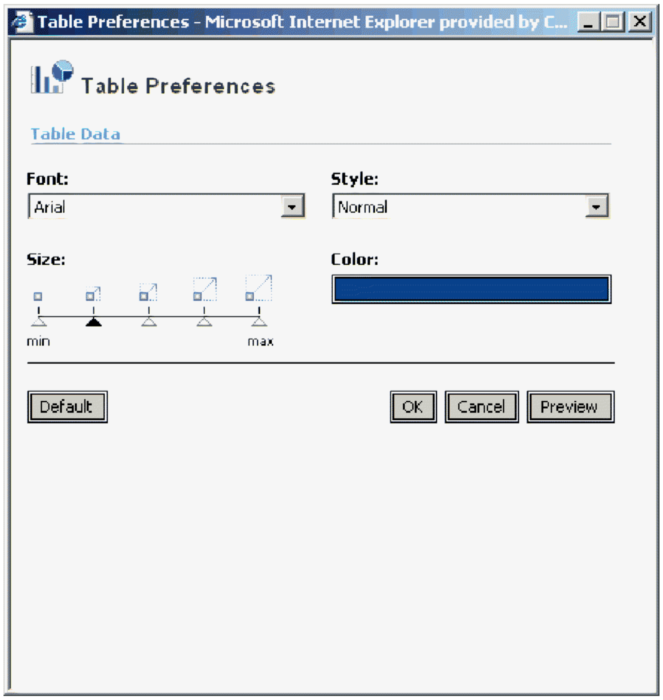

Font: Arial

Style: Normal

Figure 13. System Monitor, Table Preferences dialog box

2. In the Preferences dialog box, select the appearance options that you want to change. For a list of appearance options are available with each display option, see “Steps for using monitor display options” on page 8and Table 3 on page 9..
3. Click **Preview** to see a preview of the changes you made.
4. Click **OK**. The changes appear in the monitor. Changes to the preferences of a display option appear in all monitors that use that particular display option.

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

Steps for setting the refresh rate for run-time values

Some monitors display run-time values of a component. For these monitors, you can specify how often you want statistics to be refreshed. The refresh rate you set is for the system as a whole, not for individual components.

Perform the following steps to set the refresh rate for monitored run-time values:

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

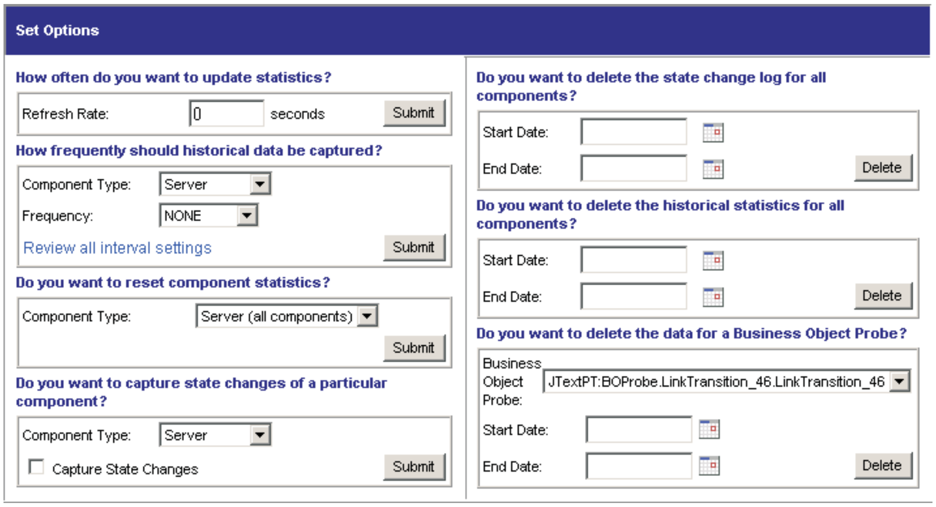

Set Options

How often do you want to update statistics?

Refresh Rate: 0 seconds

How frequently should historical data be captured?

Component Type: Server

Frequency: NONE

Review all interval settings

Do you want to reset component statistics?

Component Type: Server (all components)

Do you want to capture state changes of a particular component?

Component Type: Server

□ Capture State Changes

Do you want to delete the state change log for all components?

Start Date:

End Date:

Do you want to delete the historical statistics for all components?

Start Date:

End Date:

Do you want to delete the data for a Business Object Probe?

Business Object Probe: JTextPT:BOProbe.LinkTransition_46.LinkTransition_46

Start Date:

End Date:

Figure 14. System Monitor, Set Options dialog box

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

Steps for setting the frequency for historical data capture

Perform the following steps to set the rate at which historical data is captured:

1. Click **Set Options** in the left frame of System Monitor. The Set Options dialog box appears (see Figure 14).
2. In the **How frequently should historical data be captured?** section, click the **Review all interval settings** link. The Historical Statistics Interval Rates dialog box appears (see Figure 15 on page 20).

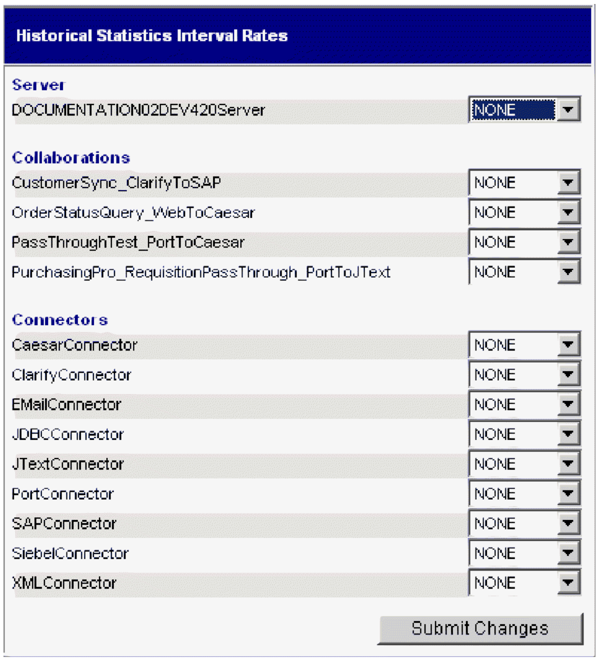

Historical Statistics Interval Rates

Server

DOCUMENTATION02DEV420Server	NONE

Collaborations

CustomerSync_ClarifyToSAP	NONE
OrderStatusQuery_WebToCaesar	NONE
PassThroughTest_PortToCaesar	NONE
PurchasingPro_RequisitionPassThrough_PortToJText	NONE

Connectors

CaesarConnector	NONE
ClarifyConnector	NONE
EMailConnector	NONE
JDBCConnector	NONE
JTextConnector	NONE
PortConnector	NONE
SAPConnector	NONE
SiebelConnector	NONE
XMLConnector	NONE

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

3. Set the interval rates for the server, for each collaboration object, and for each connector by selecting one of the following:
   - NONE
   - 15 minutes
   - 30 minutes
   - 1 hour
   - 4 hours
   - 12 hours
   - 24 hours
4. Click **Submit Changes** to submit all of the interval rates for all of the components.

Note: Alternatively, you can set the interval rate for a single component in the Set Options dialog box by selecting the component from the **Component Type** list and the interval rate from the **Frequency** list, then clicking the **Submit** button.

Steps for resetting run-time statistic values

The run-time statistics are kept in memory from the time the server is started. If the server is running for several days or weeks, these values can become very large.

Perform the following steps to reset the value of a component's run-time statistics to "0":

1. Click **Set Options** in the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 19).
2. In the **Do you want to reset component statistics?** section, select the component from the **Component Type** list.
   - If you select **Server**, then run-time statistics for all components are reset.

- If you select **Collaboration** or **Connector**, then select the component from the **Component** list. Only statistics for that component are reset.

3. Click **Submit**.

Steps for capturing state changes

Perform the following steps to configure how state changes for each component are captured and sent to a log file:

1. Click **Set Options** in the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 19).
2. Under the **Do you want to capture state changes of a particular component?** section, select the component from the **Component Type** list.

   Note: If you selected **Collaboration** or **Connector** as the component type, you are prompted to select a particular collaboration object or connector.
3. Select the **Capture State Changes** check box, then click the **Submit** button.

Steps for deleting the state change log

As the state change log grows, you may need to delete old data.

Perform the following steps to delete the log for a particular time:

1. Click **Set Options** in the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 19).
2. Under the **Do you want to delete the state change log for all components?** section, do the following:
   - Click the calendar icons to enter the start date and end date for the data to be deleted.
   - Click the **Delete** button.

Steps for deleting historical statistics

As the historical data grows, you may need to delete old data.

Perform the following steps to delete historical data for a particular time period:

1. Click **Set Options** in the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 19).
2. Under the **Do you want to delete the historical statistics for all components?** section, do the following:
   - Click the calendar icons to enter the start date and end date for the data to be deleted.
   - Click the **Delete** button.

Steps for deleting the business object probe data file

As the business object probe data grows in size, you may need to delete old data.

Perform the following steps to delete the data for a particular time period:

1. Click **Set Options** in the left frame of System Monitor. The Set Options dialog box appears (see Figure 14 on page 19).
2. Under the **Do you want to delete the data for a business object probe?** section, do the following:
   - Select the business object probe from the **Business Object Probe** list.
   - Click the calendar icons to enter the start date and end date for the data to be deleted.
   - click **Delete**.

Using persistent monitoring

Persistent monitoring is a subsystem of InterChange Server that monitors and stores historical state and statistical information of collaboration objects, connectors and the system as a whole. You can use persistent monitoring with system components or with the entire system.

You configure the various levels of persistent monitoring for system components from the Set Options dialog box in System Monitor. Those procedures are included in the previous section, "Customizing data" on page 17:

- "Steps for setting the frequency for historical data capture" on page 19
- "Steps for resetting run-time statistic values" on page 20
- "Steps for capturing state changes" on page 21
- "Steps for deleting the state change log" on page 21
- "Steps for deleting historical statistics" on page 21
- "Steps for deleting the business object probe data file" on page 21

To configure system-wide persistent monitoring, you use the Edit Configuration tool in System Manager. This section describes how to configure system-wide persistent monitoring with the Edit Configuration tool and how to access the results of system-wide persistent monitoring from System Monitor.

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

Steps for configuring system-wide parameters for persistent monitoring

Perform the following steps to configure system-wide parameters of persistent monitoring.

1. Open the Edit Configuration tool by doing the following:
   - Open System Manager.
   - Right-click the server under **Server Instances** in the InterChange Server Component Management view, and click **Edit Configuration**.

   The upper-right section of the System Manager window becomes a tool from which you can edit the InterchangeServer.cfg file.
2. Click the **Misc** tab (see Figure 16 on page 23).

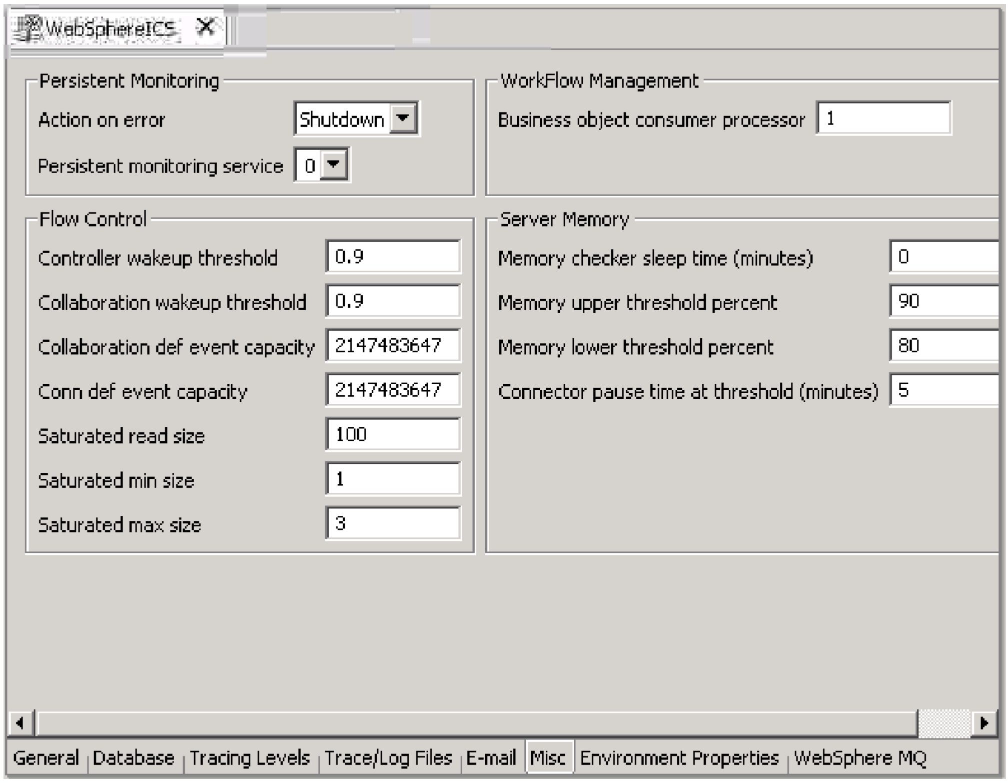

Persistent Monitoring

Action on error: Shutdown

Persistent monitoring service: 0

WorkFlow Management

Business object consumer processor: 1

Flow Control

Controller wakeup threshold: 0.9

Collaboration wakeup threshold: 0.9

Collaboration def event capacity: 2147483647

Conn def event capacity: 2147483647

Saturated read size: 100

Saturated min size: 1

Saturated max size: 3

Server Memory

Memory checker sleep time (minutes): 0

Memory upper threshold percent: 90

Memory lower threshold percent: 80

Connector pause time at threshold (minutes): 5

Figure 16. Edit Configuration tool, Misc tab

3. Under **Persistent Monitoring**, do the following:
   - If you want InterChange Server to continue running in the event of errors experienced by the persistent monitoring system, select **Continue** in the **Action on error** list.
   - If you want InterChange Server to shut down in response to errors with the subsystem, select **Shutdown** in the **Action on error** list.
   - To specify the tracing level for the subsystem, select the desired tracing level in the **Persistent monitoring service** list.

Steps for accessing the results from persistent monitoring

Perform the following steps to access the results of persistent monitoring:

1. Open System Monitor.
2. Select one of the following views in the **Views** column to display historical state and statistical information:
   - Collaboration History
   - Server History

For more information on using default views, see “Steps for using default views” on page 15. Alternatively, you can create your own views that can contain historical data. For more information on creating views, see “Steps for creating your own views” on page 16.

Using System Manager to monitor the system

You can use the InterChange Server Component Management view in System Manager to monitor the IBM WebSphere InterChange Server system and to get informational messages for all component status changes in the system. You also can start, stop, pause, and shut down InterChange Server components and change component properties from this view. For instructions on starting, stopping, and pausing components or on changing component properties, see Chapter 2, “Administering components of the system,” on page 47.

The following section describes how to work in System Manager to connect to an InterChange Server instance and to view component statistics from the InterChange Server Component Management view. To use the InterChange Server Component Management view in System Manager, you must first connect to an InterChange Server instance.

Steps for connecting to an InterChange Server instance

To use the InterChange Server Component Management view of System Manager, you must first connect to an InterChange Server instance.

Perform the following steps to connect to an InterChange Server instance:

1. From the InterChange Server Component Management view of System Manager (lower left quadrant of System Manager perspecive), right-click the InterChange Server instance you want to connect to, then select Connect.
2. Enter the following information in the Login dialog box that appears:
   - In the User Name field, type or select the user name for the server.
   - In the Password field, type the password for the server.
   - Click OK.

   When a connection is made, the light on the InterChange Server instance icon changes from red to green, and any objects that have been deployed to that server appear in folders beneath the server.

Viewing and using statistics

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

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

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

“Server statistics” on page 25

“Collaboration object statistics” on page 26

"Connector statistics" on page 27

Note: Before you can see any system statistics, System Manager must be connected to an InterChange Server instance. For instructions on connecting to an InterChange Server instance, see "Steps for connecting to an InterChange Server instance" on page 24.

Server statistics

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

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

"Database connections" on page 25

"Depth of message queues" on page 26

Database connections: Using the Database connections section of the Statistics window for InterChange Server, you can find out how many database connections the InterChange Server system's connection cache is currently using and the peak amount used since the server was booted. This can help you tune InterChange Server's interaction with the underlying Database server. By using the parameters in the `InterchangeSystem.cfg` file and the respective underlying database server's `.cfg` files, you can configure the optimal number of connections.

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

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

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

The Database connections area shows statistics for:

Cache max configured	The maximum number of connections configured. This is the value the attribute MAX_CONNECTIONS. The upper limit for this attribute is 50. If this attribute is not configured, it displays Default.
Cache in use	The current number of connections used from the connection cache.
Cache peak	The maximum number of connections used by the server from its connection cache since the server was booted.

The area below the cache statistics lists the system and dynamic connection pools. The system pools are `REPOSITORY`, `EVENTS_MANAGEMENT`, `FLOW_MONITORING`, and `TRANSACTIONS`. The dynamic pool is the Relationship pool. The following details for each of these pools are maintained:

Free
The current number of available connections in the connection pool.

In use
The current number of connections used by this connection pool.

Max Configured
The maximum number of connections configured. This is the value in the `InterchangeSystem.cfg` file for the attribute MAX_CONNECTIONS in the respective subsections of the different connection pools (Event Management, Transactions, Repository). The upper limit for this attribute is 50. If this attribute is not configured, it displays Default.

Peak
The highest number of connections used by the server from this pool since the server was booted.

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

Queue Name
The name of the subscription queue.

Current
The number of messages currently in the queue. This does not include subscriptions messages that are in the work-in-progress (WIP) queue.

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

Collaboration object statistics

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

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

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

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

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

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

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

Flow Control section: This section displays the number of buffered events and the number of events pending in the database. It also displays two configurable Flow Control properties: Max Event Capacity and Blocked Status. Use this section to monitor the Flow Control of the collaboration object and to determine if you need to reconfigure the Flow Control properties of the collaboration object. For instructions on reconfiguring Flow Control properties, see “Steps for configuring flow control for collaboration objects” on page 77.

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

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

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

Connector statistics

To check connector statistics, right-click the connector whose statistics you want to view, then select Statistics. The statistics for that connector appear in the upper-right quadrant of System Manager.

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

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

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

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

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

Flow control: This section displays the number of buffered events and the number of events pending in the database. It also displays two configurable Flow Control properties: Max Event Capacity and Blocked Status. Use this section to monitor the Flow Control of the connector and to determine if you need to reconfigure the Flow Control properties of the connector. For instructions on reconfiguring Flow Control properties, see “Steps for configuring flow control for connectors” on page 69.

Using the SNMP agent

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

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

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

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

This section covers the following topics:

“How SNMP works” on page 28

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

“What SNMP manages” on page 30

“How to use the SNMP agent” on page 31

“Using the SNMP Agent Configuration Manager” on page 33

How SNMP works

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

“SNMP architecture” on page 29

“Management Information Base (MIB)” on page 29

“Community names” on page 30

SNMP architecture

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

- SNMP manager
- SNMP agent
- Managed devices

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

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

The system devices that can be managed using the SNMP agent are InterChange Server, collaborations, and connectors. For detailed information about what types of tasks the SNMP manager can perform on each of these managed devices, see "What SNMP manages" on page 30.

Management Information Base (MIB)

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

Two MIB definitions are provided with the WebSphere InterChange Server product: `wbi_snmpagent_v2.mib` and `wbi_snmpagent_v1.mib`. These files are located in the <*WebSphere_Business_Integration_Install_Dir*>`\snmp` directory, where <*WebSphere_Business_Integration_Install_Dir*> is the directory where you installed the WebSphere InterChange Server product.

After installing and configuring the SNMP agent, import into your SNMP manager either `wbi_snmpagent_v2.mib` (if your SNMP manager supports SNMP version-2) or `wbi_snmpagent_v1.mib` (if your SNMP manager supports SNMP version-1).

For instructions on installing and configuring the SNMP agent, see the *System Installation Guide for UNIX* or *for Windows*.

For instructions on importing the MIB file to the SNMP manager, refer to the documentation provided with your SNMP management software.

Note: The MIB definitions for the SNMP agent are not configurable, but the SNMP agent MIB table is configurable. For instructions on configuring the SNMP agent MIB table, see "Using the SNMP Agent Configuration Manager" on page 33.

Community names

Access control within SNMP version-2 is supported through SNMP community names. Community names function like passwords, allowing various users to manage system components by accessing the SNMP agent using a community name. For instructions on configuring community names, see "Steps for configuring community names" on page 35.

How the SNMP agent and SNMP manager communicate

The SNMP agent and SNMP manager communicate using the Simple Network Management Protocol. Table 5 describes the requests and notifications exchanged between the SNMP agent and SNMP manager.

Table 5. Communication between SNMP agent and SNMP manager

Request or notification	Description
`Get`	The SNMP manager sends this request to the SNMP agent to get information about the device or one of its managed components.
`GetNext`	The SNMP manager sends this request to the SNMP agent to get information about the component next to the one requested previously. This is used to iterate through a table of components.
`GetBulk`	The SNMP manager sends this request to the SNMP agent to get an entire table of data.
`Set`	The SNMP manager sends this request to the SNMP agent to set a configurable parameter in the managed devices. It is also used to start and stop components.
`Trap`	A trap is an asynchronous notification sent by the SNMP agent to the SNMP manager when the status of a component in the managed device changes, and the SNMP manager has expressed interest in such status changes.

What SNMP manages

The following sections list the SNMP operations that can be performed on the managed devices by the SNMP manager. For the most up-to-date operations, refer to the current MIB file, wbi_`snmpagent_v2.mib` (for SNMP version-2) or `wbi_snmpagent_v1.mib` (for SNMP version-1), located in the `<`*`WebSphere_Business_Integration_Install_Dir`*`>\snmp` directory. This section covers the following topics:

"SNMP management of InterChange Server" on page 30

"SNMP management of collaboration objects" on page 31

"SNMP management of connectors" on page 31

SNMP management of InterChange Server

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

- Register the SNMP manager to monitor the server, and unregister it.
- Query a server's status (that is, running or stopped)
- Receive traps and log the trap name when a server's status changes
- Query for the following server parameters:
  – Time of the last boot

- Amount of time the server has been operating (uptime)
  - Number of events processed since boot time
  - Number of failed events since boot time
  - Available memory in the server (total/free)

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

SNMP management of collaboration objects

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

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

SNMP management of connectors

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

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

How to use the SNMP agent

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

- Ensure that the SNMP agent is installed and configured. For instructions on installing and configuring the SNMP agent, see the *System Installation Guide for UNIX* or *for Windows*.
- After the SNMP agent is installed and configured, you must import the correct MIB file into your SNMP manager. For SNMP managers that support SNMP version-2, import `wbi_snmpagent_v2.mib`, and for SNMP managers that support

SNMP version-1, import `wbi_snmpagent_v1.mib`. These files are located in the `<WebSphere_Business_Integration_Install_Dir>\snmp` directory.

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

This section covers the following topics:

"Steps for starting the SNMP agent" on page 32

"Steps for stopping the SNMP agent" on page 32

"Steps for reconfiguring the SNMP agent" on page 32

Steps for starting the SNMP agent

Perform the following steps to start the SNMP agent:

UNIX

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

If you run the SNMP agent on a UNIX operating system and the SNMP agent is configured to the default port (1161) or to any port number less than 1024, the port must not be in use and you must be `root` to run the SNMP startup script. If the SNMP agent is configured to a port number greater than or equal to 1024, a non-root user can start the script.

Steps for stopping the SNMP agent

Perform the following steps to stop the SNMP agent:

UNIX

Use the `snmpagent_manager` script.

Steps for reconfiguring the SNMP agent

Configuration information for the SNMP agent is stored in a configuration file named `wbi_snmpagent.cfg` in the `<WebSphere Business Integration Install Dir>\snmp\config` directory.

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

The values contained in the file are specified as:

```
ParameterName: value
```

Table 6 lists the parameters used for the operation of the SNMP agent.

Table 6. SNMP agent configuration file parameters

Parameter	Description	Values	Default
TraceLevel	Defines the verboseness of the trace information.	0-5	0

Table 6. SNMP agent configuration file parameters (continued)

Parameter	Description	Values	Default
LogFile	Path to the log file.		`wbi_snmpagent.log`, located in the <*WebSphere Business Integration Install Dir*>`\snmp\log` directory
AgentStateFile	Path to the file that contains the agent's state.		`wbi_snmpagent.sts`, located in the <*WebSphere Business Integration Install Dir*>`\snmp\state` directory.
PollInterval	Not all information required by the SNMP agent is available through callbacks, and certain information needs to be obtained through periodic polling. This parameter specifies the polling interval, in seconds	0 (no polling) and up	`30`
Port	The port on which the SNMP agent listens for requests from SNMP managers	A valid port number	`1161` (the default SNMP port number for UNIX) `161` (the default SNMP port number for Windows)

If you want to make changes to the SNMP configuration, the SNMP Configuration wizard provides fields for the information in Table 6.. The wizard creates (or modifies) the `wbi_snmpagent.cfg` file based on the values in these fields.

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

Using the SNMP Agent Configuration Manager

The SNMP Agent Configuration Manager allows you to configure the MIB tables associated with the SNMP agent. For more information about MIB tables, see "Management Information Base (MIB)" on page 29. These tables include the Community table, the Trap Forwarding table, and the Server Access table.

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

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

This section covers the following topics:

"Installing the SNMP Agent Configuration Manager" on page 34

"Steps for starting the SNMP Agent Configuration Manager" on page 34

"Steps for connecting to the SNMP agent" on page 34

"Steps for configuring community names" on page 35

"Steps for configuring trap forwarding entries" on page 36

"Steps for configuring server access entries" on page 38

"Steps for configuring RBAC security" on page 39

Installing the SNMP Agent Configuration Manager

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

Steps for starting the SNMP Agent Configuration Manager

Perform the following steps to start the SNMP Agent Configuration Manager, depending on your operating system:

UNIX

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

The SNMP Agent Configuration Manager appears.

Before you can begin editing the MIB table, you must connect to the SNMP agent. See "Steps for connecting to the SNMP agent" for instructions.

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

Steps for connecting to the SNMP agent

Perform the following steps to connect the SNMP Agent Configuration Manager to a running SNMP agent:

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

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

Steps for configuring community names

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

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

Adding community names: You add community names to the MIB table when you want to give new users permission to manage system components.

Perform the following steps to add a community name:

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

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

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

Perform the following steps to edit a community name:

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

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

Removing community names: You remove community names when you want to completely delete them from the MIB table.

Perform the following steps to remove a community name:

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

Steps for configuring trap forwarding entries

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

When a trap is sent, the SNMP Agent notifies the designated host:port specified in the Trap Forwarding Table and logs the trap name, for example, collabTrapEventsLongTime, as well as the connector application name, connector server name and connector application statur. A monitoring network manager receives the trap, which triggers a response, for example, sending email to the System Administrator.

The Trap Forwarding table has four entries:

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

Table 7 list the SNMP Traps that exist in InterChange Server.

Table 7. SNMP Traps

Trap Type	Variables	Description	Trap ID
serverTrapStatus	serverName, serverStatus	Generated when a server starts or stops	1
collabTrapStatus	collabName, collabServerName, collabStatus	Generated when a collaboration starts or stops	2
collabTrapEventsFailed	collabName, collabServerName, collabEventsTrgdFlwFailed	Generated when event fails on collaboration	3
collabTrapEventsLongTime	collabName, collabServerName	Generated when an event takes longer than a specified time on collaboration	4
connTrapAgentStatus	connName, connServerName, connAgentStatus	Generated when a connector agent's status is changed	5
connTrapStatus	connName, connServerName, connStatus	Generated when a connector status is changed	6
connTrapAppStatus	connName, connServerName, connAppStatus	Generated when a connector application's status is changed	7

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

Perform the following steps to add a trap forwarding entry to the MIB table:

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

   Note: Each Trap ID must be unique in the network.
3. In the Manager Host field, type the host name or IP address of the machine where the SNMP manager runs.

   Note: For UNIX users who use the host name, be sure to use the correct case.
4. In the Trap Port field, type the port number that the SNMP manager uses to listen to traps.
5. The Row Status drop-down menu is not configurable when adding a trap forwarding entry. The default setting is ″create & go.″
6. Click OK. The new trap forwarding entry appears in the Trap Forwarding table with the Row Status set to Active.

Editing trap forwarding entries: You edit trap forwarding entries when you want to change the Manager Host, Trap Port, or Row Status information of registered trap forwarding entries.

Perform the following steps to edit an existing trap forwarding entry:

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

   Note: You cannot edit the Trap ID.
2. Change the Manager Host by typing a different host name or IP address of the machine where the SNMP manager runs.

   Note: For UNIX users who use the host name, be sure to use the correct case.
3. Change the Trap Port by typing a different port number that the SNMP manager uses to listen to traps.
4. Change the Row Status of the trap forwarding entry by selecting one of the following options from the Row Status drop-down menu:
   - **active**: This activates the trap forwarding entry, allowing it to be used to access the SNMP agent.
   - **not in service**: This deactivates the trap forwarding entry, but it stores the trap forwarding information in the table. Choose this option if you know you want to reactivate this trap forwarding entry at some point in the future.
   - **destroy**: This removes the chosen community name entry from the MIB table.
5. Click OK. Any changes you made appear in the Trap Forwarding table.

Removing trap forwarding entries: You remove trap forwarding entries when you want to completely remove them from the MIB table.

Perform the following steps to remove a trap forwarding entry:

1. From the Trap Forwarding tab, select the Trap ID you want to remove.

2. Click Remove. The trap forwarding entry is removed from the Trap Forwarding table.

Steps for configuring server access entries

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

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

Adding server access entries: Perform the following steps to create a new server access:

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

   Note: For UNIX users who use the host name, be sure to use the correct case.
3. In the WebSphere InterChange Server field, type the InterChange Server name or IP address.

   Note: For UNIX users who use the host name, be sure to use the correct case.
4. The Row Status drop-down menu is not configurable when adding a Server Access entry. It is set to ″create & go.″
5. Click OK. The new server access entry is added to the Server Access table, with the Row Status set to ″not in service.″

Editing server access entries: You edit the server access entries of the MIB table when you want to change the row status of the server access entries.

Perform the following steps to edit a server access entry, do the following:

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

     Note: In order to change the status from ″not in service″ to ″active,″ both the SNMP manager and the InterChange Server being managed must be running.
   - **not in service**: This deactivates the manager host, but it stores the manager host information in the table. Choose this option if you know you want to reactivate this manager host at some point in the future.
   - **destroy**: This removes the chosen manager host from the MIB table.
3. Click OK. Any changes you made appear in the Server Access table.

Removing server access entries: You remove server access entries when you want to completely delete them from the MIB table.

Perform the following steps to remove a server access entry:

1. From the Server Access tab, select the server access entry you want to remove.

2. Click Remove. The Server Access entry is removed from the Server Access table.

Steps for configuring RBAC security

Role-based access control (RBAC) supports multiple users and enhanced security features based on roles. A role is a collection of users who share common functionality. Assigning functions into roles allows the administrator to work more effectively by reducing the burden on the administrator during the assignment of permissions.

Due to the addition of RBAC functionality, the SNMP agent now allows the input of usernames and passwords to help administer these roles. If RBAC security is enabled on InterChange Server, a user must specify a username and password to connect to the InterChange Server.

The RBAC Security table has four entries:

- **WebSphere ICS**: The unique ID for the WebSphere InterChange Server
- **User Name**: The username assigned to a specific individual
- **Password**: The password assigned to the username
- **Row Status**: The status of the trap (active or not in service)

Adding RBAC security entries: You add RBAC security entries to the MIB table to connect the SNMP agent to the InterChange Server when RBAC security is enabled.

Perform the following steps to add an RBAC security entry to the MIB table:

1. From the RBAC Security tab, click Add. The RBAC Security Table Item dialog box appears.
2. In the WebSphere ICS field, type an integer.
3. In the Username field, type the username assigned to the role.
4. In the Password field, type the password assigned to the username.
5. The Row Status drop-down menu is not configurable when adding an RBAC security entry. It is set to ″create & go.″
6. Click OK. The new RBAC security entry appears in the RBAC security table with the Row Status set to Active.

Editing RBAC security entries: You edit RBAC security entries when you want to change the Username, Password or Row Status information of registered RBAC security entries.

Perform the following steps to edit an existing RBAC security entry:

1. In the RBAC Security tab, select a Username, then click Edit. The RBAC Security Table Item dialog box appears.
2. Change the Username by typing a different username in the available space.
3. Change the Password by typing a different password in the available space.
4. Change the Row Status of the RBAC security entry by selecting one of the following options from the Row Status drop-down menu:
   - **active**: This activates the RBAC security entry, allowing it to be used to access the SNMP agent.
   - **not in service**: This deactivates the RBAC security entry, but it stores the RBAC security information in the table. Choose this option if you know you want to reactivate this RBAC security entry at some point in the future.

- **destroy**: This removes the chosen RBAC security entry from the MIB table.

5. Click OK. Any changes you made appear in the RBAC Security table.

Removing trap forwarding entries: You remove RBAC security entries when you want to completely remove them from the MIB table.

Perform the following steps to remove a RBAC security entry:

1. From the RBAC Security tab, select the Trap ID you want to remove.
2. Click Remove. The RBAC security entry is removed from the RBAC Security table.

Using WebSphere Business Integration Monitor

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

- WebSphere Business Integration Monitor, version 4.2.4, Fix pack 1, must be installed on the machine from where you intend to do the monitoring.
- MQ Workflow adapter must be installed and enabled with **flow monitoring**. Flow monitoring is a service that provides a view of the event information associated with each flow as it passes through a collaboration in InterChange Server.
- InterChange Server, version 4.2.2.2, must be configured for flow monitoring.
- Enable flow monitoring for the collaboration with the Auditing Collaboration Utility provided by the WebSphere Business Integration Monitor. For further information, refer to the Administration Guide in the WebSphere Business Integration Monitor documentation.

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

Important: WebSphere Business Integration Monitor will not work with InterChange Server if InterChange Server if it is using Microsoft SQL Server as its database.

This section covers the following topics:

"How WebSphere Business Integration Monitor works with InterChange Server" on page 41

"Steps for enabling flow monitoring in the MQWorkflow adapter" on page 41

"Configuring InterChange Server for flow monitoring" on page 41

"Steps for configuring tracing for flow monitoring" on page 46

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

How WebSphere Business Integration Monitor works with InterChange Server

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

Note: WebSphere Business Integration Monitor monitors InterChange Server events only, not the business data associated with those events.

Steps for enabling flow monitoring in the MQWorkflow adapter

WebSphere Business Integration Monitor can monitor an InterChange Server system only when the MQWorkflow adapter is enabled for flow monitoring. Flow monitoring is a service that provides a view of the event information associated with each flow as it passes through a collaboration in InterChange Server.

Note: You must configure WebSphere MQ to process BiDi data in order to enable communication with the BiDi environment. For more information, please refer to *System Installation Guide for Windows.*

Perform the following steps to enable the MQWorkflow adapter with flow monitoring.

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

Configuring InterChange Server for flow monitoring

WebSphere Business Integration Monitor can monitor an InterChange Server system only when that InterChange Server system is first configured with flow monitoring. Flow monitoring is a service that provides a view of the event information associated with each flow as it passes through a collaboration in InterChange Server. The following sections describe how to configure InterChange Server for flow monitoring.

“Steps for configuring the InterChange Server database for flow monitoring”

“Steps for configuring maximum queue depth for flow monitoring” on page 45

“Steps for configuring tracing for flow monitoring” on page 46

Steps for configuring the InterChange Server database for flow monitoring

The InterChange Server database can be configured for flow monitoring only after the InterChange Server installation is complete. Two tools allow you to configure the InterChange Server database for flow monitoring: the InterChange Server Configuration Wizard and System Manager.

Perform the following steps to configure the InterChange Server database for flow monitoring, using either of these tools.

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

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

   InterChange Server Configuration Wizard

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

   The InterChange Server Configuration Wizard opens.

   System Manager

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

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

2. Select the Database tab (of either the InterChange Server Configuration Wizard or the Edit Configuration window of System Manager), then scroll to the bottom of the window to see the Flow Monitoring section (see Figure 17 on page 43 for the InterChange Server Configuration Wizard orFigure 18 on page 44 for the Edit Configuration window of System Manager).

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

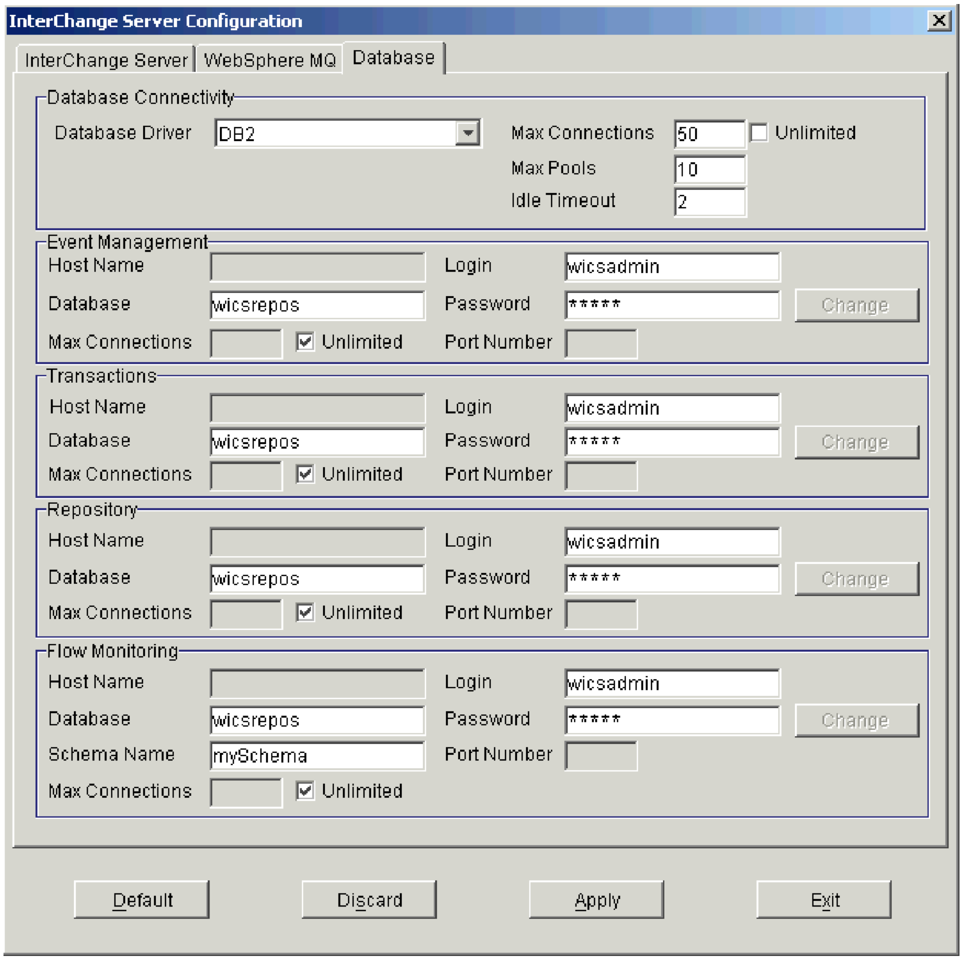

InterChange Server Configuration

Database Connectivity

Database Driver: DB2

Max Connections: 50 □ Unlimited

Max Pools: 10

Idle Timeout: 2

Event Management

Host Name:

Login: wicsadmin

Database: wicsrepos

Password: *****

Max Connections: ☑ Unlimited

Port Number:

Transactions

Host Name:

Login: wicsadmin

Database: wicsrepos

Password: *****

Max Connections: ☑ Unlimited

Port Number:

Repository

Host Name:

Login: wicsadmin

Database: wicsrepos

Password: *****

Max Connections: ☑ Unlimited

Port Number:

Flow Monitoring

Host Name:

Login: wicsadmin

Database: wicsrepos

Password: *****

Schema Name: mySchema

Port Number:

Max Connections: ☑ Unlimited

Figure 17. Database tab of the InterChange Server Configuration Wizard

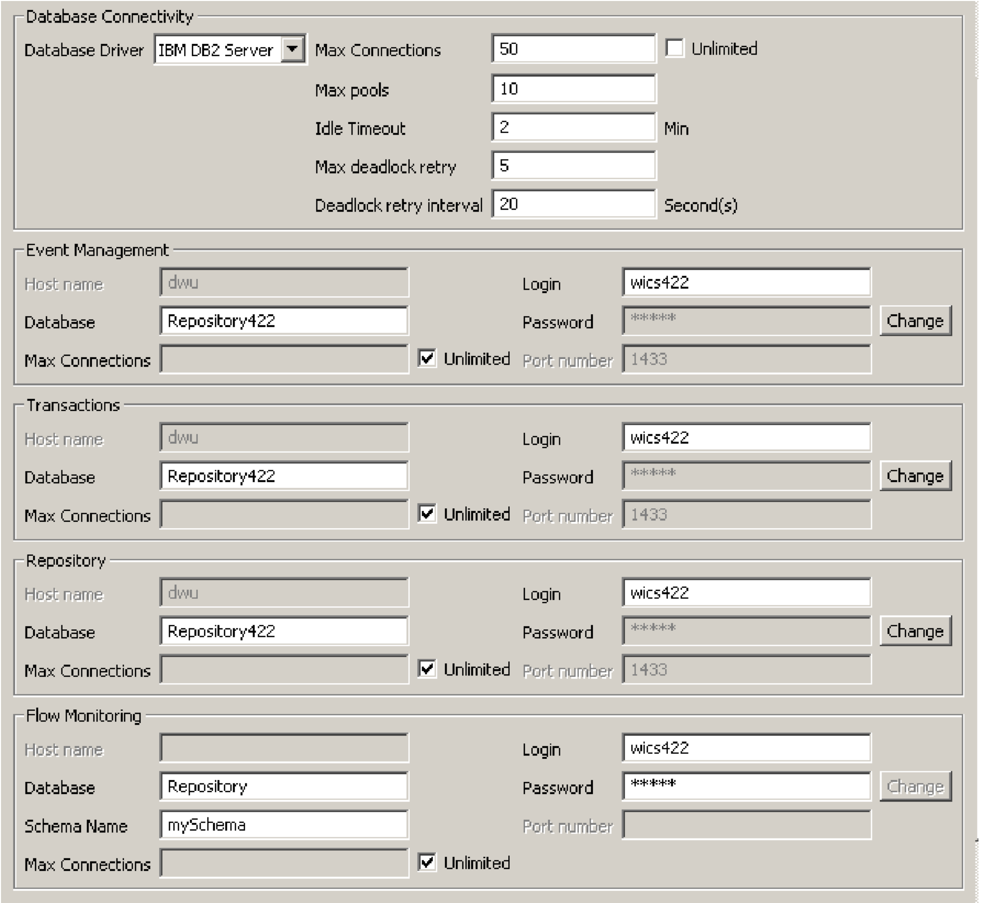

Database Connectivity

Database Driver: IBM DB2 Server

Max Connections: 50 □ Unlimited

Max pools: 10

Idle Timeout: 2 Min

Max deadlock retry: 5

Deadlock retry interval: 20 Second(s)

Event Management

Host name: dwu

Database: Repository422

Max Connections: ☑ Unlimited

Login: wics422

Password: ***** Change

Port number: 1433

Transactions

Host name: dwu

Database: Repository422

Max Connections: ☑ Unlimited

Login: wics422

Password: ***** Change

Port number: 1433

Repository

Host name: dwu

Database: Repository422

Max Connections: ☑ Unlimited

Login: wics422

Password: ***** Change

Port number: 1433

Flow Monitoring

Host name:

Database: Repository

Schema Name: mySchema

Max Connections: ☑ Unlimited

Login: wics422

Password: ***** Change

Port number:

Figure 18. Database tab of the Edit Configuration window of System Manager

3. Enter information in the Flow Monitoring fields, using the information in Table 8 as a guide.

Table 8. Flow Monitoring configuration information

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

Table 8. Flow Monitoring configuration information (continued)

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

Steps for configuring maximum queue depth for flow monitoring

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

Perform the following steps to configure the maximum queue depth from System Manager.

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

1. From the InterChange Component Management view, right-click the InterChange Server instance whose maximum queue depth you want to configure, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the `InterchangeSystem.cfg` file.
2. Click the Misc tab (see Figure 19).

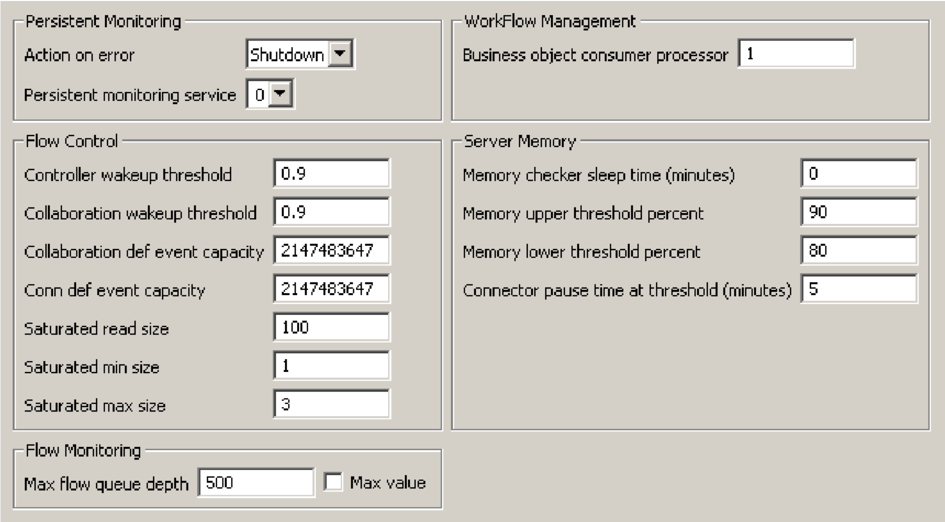

Persistent Monitoring
Action on error: Shutdown
Persistent monitoring service: 0

WorkFlow Management
Business object consumer processor: 1

Flow Control
Controller wakeup threshold: 0.9
Collaboration wakeup threshold: 0.9
Collaboration def event capacity: 2147483647
Conn def event capacity: 2147483647
Saturated read size: 100
Saturated min size: 1
Saturated max size: 3

Server Memory
Memory checker sleep time (minutes): 0
Memory upper threshold percent: 90
Memory lower threshold percent: 80
Connector pause time at threshold (minutes): 5

Flow Monitoring
Max flow queue depth: 500 ☐ Max value

Figure 19. Mics tab of InterChange Server Edit Configuration window

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

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

Steps for configuring tracing for flow monitoring

When you configure flow monitoring for an InterChange Server system, you can set a desired trace level for traced events.

Perform the following steps to set the trace level for flow monitoring.

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

1. From the InterChange Component Management view, right-click the InterChange Server instance whose flow monitoring trace level you want to configure, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the `InterchangeSystem.cfg` file.
2. Select the Tracing Levels tab.
3. In the IBM WebSphere Business Integration System Trace Levels group box, look for the "Trace level" field for Flow Monitoring.
4. In the Flow Monitoring Trace level field, select a number from 0-5 to represent the desired trace level. Table 9 describes the different trace levels for flow monitoring.

Table 9. Flow monitoring trace levels

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

Chapter 2. Administering components of the system

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

"Overview of administering the system"

"Administering InterChange Server" on page 48

"Using System Manager" on page 53

"Administering connectors" on page 57

"Administering JMS transport optimization" on page 70

"Administering collaboration objects" on page 72

"Administering maps" on page 78

"Administering relationships" on page 81

"Using Relationship Manager" on page 85

"Backing up system components" on page 104

"Using repos_copy" on page 108

"Scheduling jobs" on page 99

"Administering end-to-end privacy" on page 119

"Administering role-based access control (RBAC)" on page 123

Overview of administering the system

To begin administering the system, you must start all necessary components of the IBM WebSphere InterChange Server.

The recommended order for starting up the system is: IBM WebSphere MQ Listener, InterChange Server, then IBM WebSphere System Manager. Connectors are automatically started when you start InterChange Server.

There is some flexibility for starting components. MQ Listener can be started later, however, the connectors that depend on it are started in a paused state. System Manager can be open at any time, but you must connect to the server instance again after starting the WebSphere InterChange Server. However, if you are using IBM Java Object Request Broker (ORB) on a different machine than WebSphere InterChange Server, you must start the ORB before starting WebSphere InterChange Server.

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

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

The following tasks describe the recommended order to start the system:

1. Check that all necessary third-party software is running.
2. If you are using IBM Java Object Request Broker (ORB) on a different system than where WebSphere InterChange Server is located, start the ORB. For more information about using ORB, see “Administering the Object Request Broker” on page 150.
3. If you are using IBM WebSphere MQ, start the MQ Listener.
4. Start InterChange Server. See “Steps for starting InterChange Server” on page 48.
5. Start any connectors not automatically started. See “Starting, stopping, and pausing connectors” on page 59.
6. Start collaborations. See “Viewing collaboration object states” on page 72.
7. Start System Manager. See “Using System Manager” on page 53.

Administering InterChange Server

Administering InterChange Server may involve starting and stopping the system and managing the startup parameters and database passwords. This section includes the following topics:

“Steps for starting InterChange Server”

“Steps for shutting down InterChange Server” on page 50

“Changing the InterChange Server and database passwords” on page 51

Steps for starting InterChange Server

Perform the following steps to start InterChange Server:

UNIX

Run the `ics_manager -start` script.

Note: If the ICS password is changed, you must use the following script:
`-start -u`*`loginName`* `-p`*`password`*

Windows

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

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

InterChange Server startup parameters

Perform the following steps to customize the InterChange Server startup parameters:

UNIX

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

```
-start

-stop

-status
```

Windows

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

The parameters in table Table 10 customize the startup of InterChange Server.

Table 10. InterChange Server startup parameters

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

Note: Usage of the parameters `-u` and `-p` is not recommended as the password can be seen in clear text. It is recommended to specify these values using the `startServerUserName` and `startServerPassword` fields on the Security-RBAC tab. For more information on the Security-RBAC tab, see “Administering role-based access control (RBAC)” on page 123.

Steps for shutting down InterChange Server

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

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

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

UNIX

Run the `ics_manager -stop` script.

Note: If the ICS password is changed, you must use the following script: `-stop-u`*`loginName`* `-p`*`password`*

Windows

In InterChange Server Component Management view of System Manager, right-click the ICS instance, then select Shut Down > Gracefully. Alternatively, you can select Shut Down > Immediately, which shuts down the server without cleanup. To determine which type of shutdown is best for you, refer to the following topics: "Graceful shutdown" on page 50 and "Immediate shutdown" on page 51.

Note: A high-availability (HA) system cannot be shut down using System Monitor, SNMP manager, or the InterChange Server Component Management view of System Manager. Shutting down the server using any of these tools causes the HA system to perform a failover. To shut down the HA system, use the MSCS Administrator and use the Take Offline context menu option for all connectors and then InterChange Server. For instructions, see "Changing the status of a resource" on page 150. For more information about managing an HA system, see "Administering High-Availability (HA) systems" on page 148..

Graceful shutdown

Gracefully shutting down the system allows all currently processing and queued flows to complete before shutting down. This method may be time consuming, since all flows waiting to be processed by a running collaboration must complete prior to shutdown. However, no new flows are accepted.

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

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

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

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

Note: This procedure cannot be used to shut down ahigh-availability (HA) system. For information on shutting down an HA system, see “Changing the status of a resource” on page 150..

Immediate shutdown

Immediately stopping the system forces the system to shut down without processing any more flows. Running connectors and collaborations are stopped immediately. When the system is restarted, flows that were interrupted by the immediate shutdown are redelivered in the same processing order. If one of these flows wrote data to an application, when the flow is redelivered, it tries to duplicate the data and fails because the data already exists. If the collaboration processing the flow is transactional, a rollback occurs. If the flow is not transactional, it is moved to the resubmission queue. See “Administering failed events” on page 131 for more information on submitting a flow that fails to process.

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

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

Changing the InterChange Server and database passwords

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

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

The InterChange Server user name and password are required during repository copy and restoration when the `repos_copy` command is used. See “Using repos_copy” on page 108.

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

“Steps for changing the InterChange Server password”

“Steps for changing the database passwords” on page 52

Steps for changing the InterChange Server password

To change the password for InterChange Server:

1. Open System Manager.
2. Right-click the InterChange Server instance in the InterChange Server Component Management view, then select **Change Password**. The Change InterChange Server Password dialog box appears.

3. Enter the current password in the **Old Password** field.
4. Enter a new password in the **New Password** field.
5. Reenter the new password in the **Confirm Password** field.
6. Click **OK**.

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

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

Steps for changing the database passwords

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

To change the database passwords:

1. In the InterChange Server Component Management view, right-click the InterChange Server instance whose database password you want to change, then select **Edit Configuration**. The upper-right section of the window changes to an editing tool in which many system properties can be changed.
2. Click the Database tab to access the database configuration properties. The Server Property and Configuration window for database properties appears (see Figure 20).

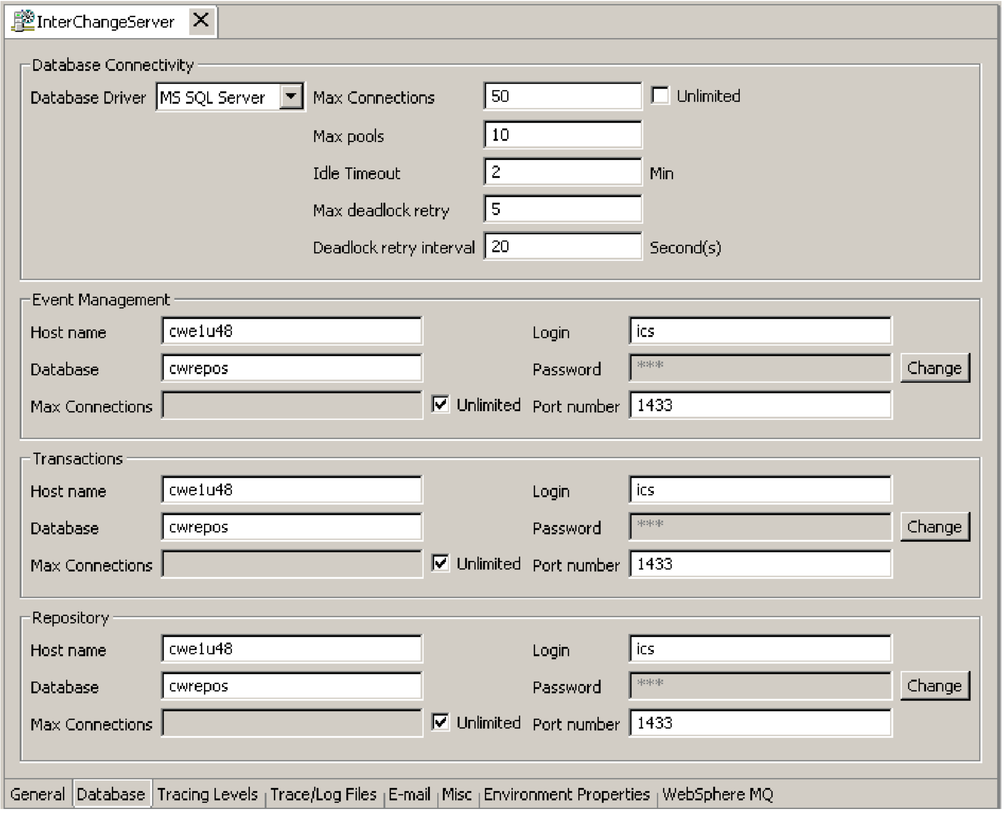

InterChangeServer

Database Connectivity

Database Driver: MS SQL Server

Max Connections: 50 □ Unlimited

Max pools: 10

Idle Timeout: 2 Min

Max deadlock retry: 5

Deadlock retry interval: 20 Second(s)

Event Management

Host name: cwe1u48

Login: ics

Database: cwrepos

Password: ***

Max Connections: ☑ Unlimited

Port number: 1433

Transactions

Host name: cwe1u48

Login: ics

Database: cwrepos

Password: ***

Max Connections: ☑ Unlimited

Port number: 1433

Repository

Host name: cwe1u48

Login: ics

Database: cwrepos

Password: ***

Max Connections: ☑ Unlimited

Port number: 1433

Figure 20. Database tab of Edit Configuration window

3. Change any of the database passwords:

a. In the section for the appropriate database (Event Management, Transactions, or Repository), click the **Change** button.

   A dialog box for changing the password appears.

b. Type the old password in the **Old Password** field.

c. Type a new password in the **New Password** field.

   A maximum of 30 characters is allowed.

d. Retype the new password in the **Confirm Password** field.

4. Click **OK**.

Using System Manager

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

"Steps for starting System Manager"

"Steps for shutting down System Manager" on page 54

"Steps for refreshing System Manager and updating components" on page 54

"Steps for configuring system-wide flow control" on page 55

Steps for starting System Manager

Perform the following step to start System Manager:

Click **Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration Toolset > Administrative > System Manager**.

The System Manager perspective of the IBM WebSphere Studio Workbench appears (see Figure 21 on page 54).

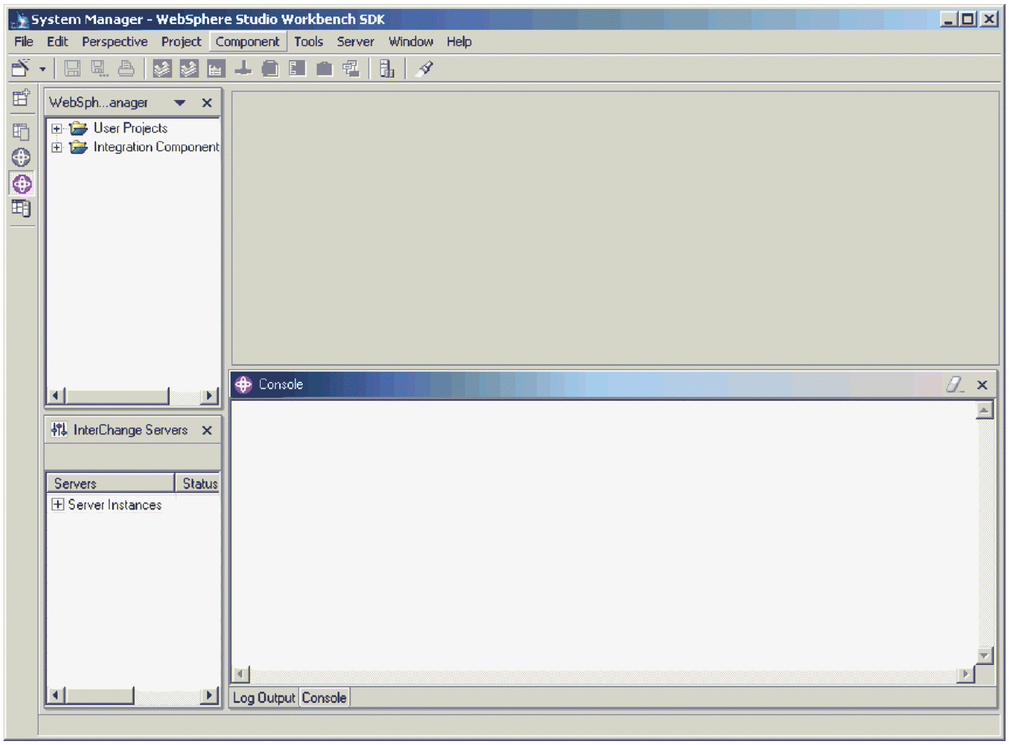

Figure 21. System Manager

Steps for shutting down System Manager

Perform the followin steps to shut down System Manager:

In IBM WebSphere Studio Workbench, select **File > Exit**.

Note: Be sure to shut down any InterChange Servers before shutting down System Manager. For instructions on shutting down InterChange Servers, see "Steps for shutting down InterChange Server" on page 50.

Steps for refreshing System Manager and updating components

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

Perform the following step to refresh InterChange Server:

Right-click the server under Server Instances, then select **Refresh**.

The following describes which components can be updated during system run time:

Business objects
Not updated during run time. The repository is read only once when InterChange Server starts up.

Collaboration object properties
Updated during run time. For example, collaboration object trace levels take effect as soon as they are set.

Collaboration object code changes
Updated during system run time.

Map code changes
Updated during system run time. If mapping code is updated and recompiled, connectors must be rebound to the altered maps.

Steps for configuring system-wide flow control

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control for individual components, see “Steps for configuring flow control for connectors” on page 69 or “Steps for configuring flow control for collaboration objects” on page 77.

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

To monitor how flow control is working in the system, you can view the Flow Control monitor and view provided as part of System Monitor or you can view the Statistics for collaboration objects or connectors from the InterChange Server Component Management view of System Manager. For more information on using the Flow Control monitor and view in System Monitor, see “Steps for reviewing default monitors” on page 2 and “Steps for using default views” on page 15. For more information on viewing the flow control from the InterChange Server Component Management view of System Manager, see “Collaboration object statistics” on page 26 or “Connector statistics” on page 27.

Perform the following steps to configure system-wide flow control:

1. In the InterChange Server Component Management view of System Manager, right-click the InterChange Server instance for which you want to configure flow control, then select **Edit Configuration**. The upper-right quadrant of System Manager changes to an editing tool in which many system properties can be changed.
2. Click the **Misc** tab. A dialog box appears with a Flow Control section (see Figure 22 on page 56).

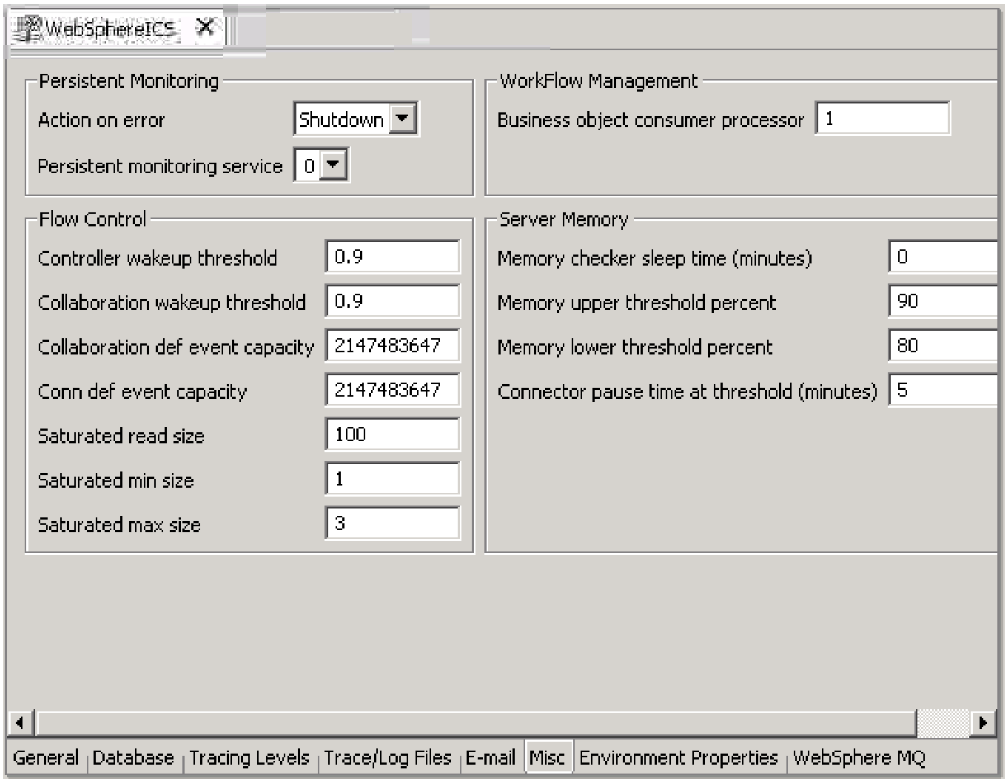

WebSphereICS

Persistent Monitoring

Action on error: Shutdown

Persistent monitoring service: 0

WorkFlow Management

Business object consumer processor: 1

Flow Control

Controller wakeup threshold: 0.9

Collaboration wakeup threshold: 0.9

Collaboration def event capacity: 2147483647

Conn def event capacity: 2147483647

Saturated read size: 100

Saturated min size: 1

Saturated max size: 3

Server Memory

Memory checker sleep time (minutes): 0

Memory upper threshold percent: 90

Memory lower threshold percent: 80

Connector pause time at threshold (minutes): 5

Figure 22. Edit configuration tool, Misc tab

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

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

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

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

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

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

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

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

4. Select **File > Save <*ServerName*>** to save the changes you made to the InterChange Server configuration.
5. Restart InterChange Server.

Administering connectors

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

While administering connectors, you can also decide whether to optimize your JMS Transport. For additional information, please refer to "Administering JMS transport optimization" on page 70.

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

This section covers the following topics:

"Viewing connector states" on page 57

"Starting, stopping, and pausing connectors" on page 59

"Steps for configuring flow control for connectors" on page 69

Viewing connector states

You can view the state of a connector either by logging on to System Monitor and opening a view that contains connector states or by using the InterChange Server Component Management view of System Manager. To log on to System Monitor, follow the instructions in "Steps for logging on to System Monitor" on page 12. To use the InterChange Server Component Management view of System Manager, follow the instructions in "Steps for connecting to an InterChange Server instance" on page 24..

The state of a connector is represented differently, depending on which tool you are using either System Monitor or System Manager.

Steps for using System Monitor to view connector states

Perform the following steps to see the state of connectors using System Monitor:

1. If the System Overview view is not displayed, click the System Overview link under Views in the left pane of the Web page. The System Overview Monitor appears (see Figure 11 on page 13) in the body of the Web page.

When the product is installed, the default view is set to System Overview, and the default monitor contained in that view is set to System Overview. These defaults can be changed to suit your monitoring needs. See “Setting up views to monitor the system” on page 14 for instructions.

2. Click the triangle next to the name of the server to reveal a list of components on the system.
3. Click the triangle next to a running collaboration to reveal its associated connectors (see Figure 23).

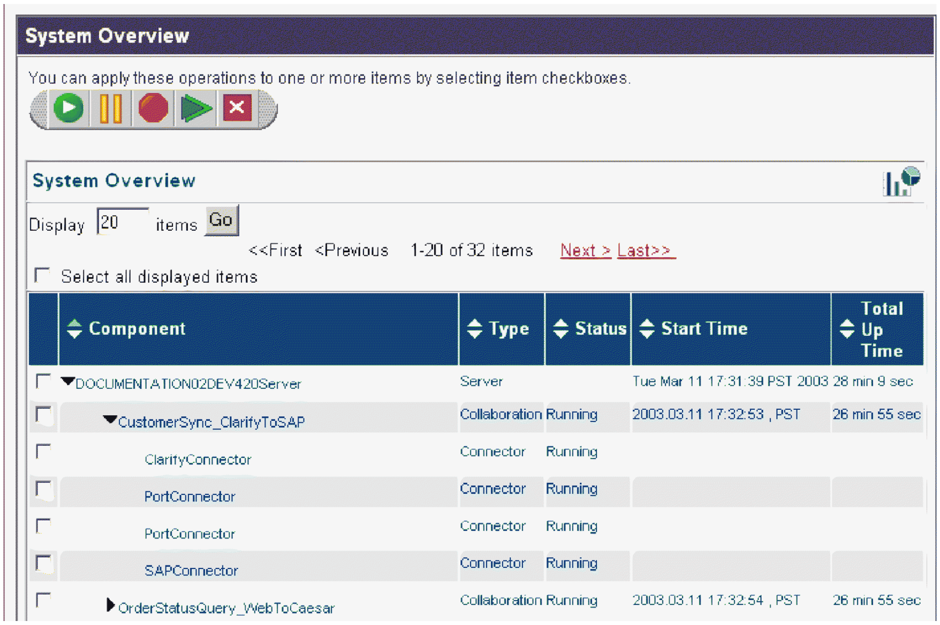

System Overview

You can apply these operations to one or more items by selecting item checkboxes.

System Overview

Display 20 items Go

<<First <Previous 1-20 of 32 items Next > Last>>

□ Select all displayed items

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

Figure 23. System Monitor, System Overview displaying connector status

Note: You may also view connector states using the Connector Overview view.

Steps for using System Manager to view connector states

Perform the following steps to view the state of a connector in System Manager:

1. Connect to the InterChange Server instance that contains the connector you want to view. See “Steps for connecting to an InterChange Server instance” on page 24 for instructions on connecting to an InterChange Server instance.
2. Expand the InterChange Server instance, then expand the Connectors folder.

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

Table 11 lists the connector states represented by the display color of each connector and shows what actions are being performed during that state.

Table 11. Connector States

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

Starting, stopping, and pausing connectors

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

“Connector initialization” on page 59

“Steps for starting, stopping, and pausing connectors using System Monitor” on page 59

“System Manager commands for changing connector states” on page 60

“Steps for manually starting a connector” on page 60

“Shutting down a connector” on page 65

“Restarting a connector” on page 66

“Steps for setting automatic and remote restart for a connector” on page 67

Connector initialization

The first time you start a connector, it must be initialized. Initializing a connector requires that you start it manually. For instructions on manually starting a connector, see “Steps for starting a connector manually on UNIX” on page 60 or “Steps for starting a connector manually on Windows” on page 63.

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

After the connector has been initialized, you can start, stop, and pause it using System Monitor or System Manager.

Steps for starting, stopping, and pausing connectors using System Monitor

Perform the following steps to start, stop and pause connectors using System Monitor:

1. While viewing the System Overview view (see Figure 23 on page 58), select a connector by placing check in the box to its left.
2. Select the Start, Pause, or Stop icon from the icon group in the upper-left corner of the view (see Figure 24).

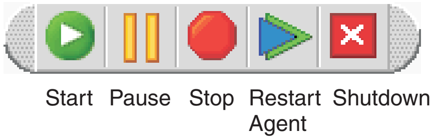

Figure 24. System Monitor, icons for starting, pausing, restarting, or shutting down components

Note: You may also start, stop or pause connectors from the Connector Overview. Connector Agents can be shutdown or restarted from this same view.

Steps for starting, stopping, and pausing connectors using System Manager

Perform the following steps to start, stop and pause connectors using System Manager:

1. From the Connectors folder in the InterChange Server Component Management view of System Manager (see Figure 29 on page 79), right-click the name of a connector.
2. Select one of the start, pause, or stop connector options.

System Manager commands for changing connector states

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

Start *Name*Connector
: Starts the connector for the application *Name* if it is paused or stopped. Connectors poll the application and connector controllers read the persistent queue. Flows are processed.

Pause *Name*Connector
: Pauses the connector for the application *Name* if it is running or stopped. Connectors stop polling the application and connector controllers stop reading the persistent queue. Flows are not processed.

Stop *Name*Connector
: Stops the connector for the application *Name* if it is running or paused. Connectors stop polling the application and fail requests with an exception message. Connector controllers stop reading the persistent queue. Flows and requests are not processed.

Shut Down *Name*Connector
: Shuts down the connector for the application *Name*. The connector's process is stopped.

Boot Up Connector Agent
: Restarts the connector for the application *Name*. This action is available only if you have set the OADAutoRestartAgent property of the connector to True. See "Steps for setting automatic and remote restart for a connector" on page 67.

Steps for manually starting a connector

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

- "Steps for starting a connector manually on UNIX"
- "Steps for starting a connector manually on Windows" on page 63

Steps for starting a connector manually on UNIX: To start a connector, use the connector manager script in the *AdapterFrameworkProductDir*/bin directory with the following syntax:

```
connector_manager_connector -start
```

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

```
AdapterFrameworkProductDir/connectors
```

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

```
connector_manager_EMail -start
```

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

The `connector_manager_connector` script is a wrapper for the generic connector manager script (`AdapterFrameworkProductDir/bin/connector_manager`). This wrapper includes the following information so that you do not need to specify it:

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

  For example:

  – The SAP connector requires the `-t` command-line option. Therefore, its startup script already includes the `-t` option.

  – All UNIX connectors are started with the `-b` option. Therefore, all connector startup scripts already include the `-b` option. To have a connector run in the foreground, remove the `-b` option from the generic connector manager script (`connector_manager`).
- The name and path of the configuration file

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

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

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

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

  – A product Installer version that assumes that you are using InterChange Server. If you used this installer, it stored configuration information for the connector in the InterChange Server repository, and it also generated a local configuration file for the connector, and set the file name and path value for that configuration file in the `AGENTCONFIG_FILE` variable in the `connector_manager_connector` script. The connector will look for a configuration file that has the file name and path specified in the

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

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

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

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

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

The generic connector manager script calls the appropriate start_*connector*.sh script, which handles the actual connector management for the connector. The IBM WebSphere Business Integration Adapter product provides a start_*connector*.sh script with each connector. This start_*connector*.sh script supports the options in Table 12 for starting a connector.

Table 12. Command-Line Options for the start_connector.sh Script

Option	Additional Information
-b	This option runs the connector as a background thread; that is, no input is read from STDIN (standard input). The generic connector_manager script (called by each connector_manager_*connector* script) automatically specifies the -b option when it invokes the start_*connector*.sh script for a connector. You can remove this option from the start_*connector*.sh invocation to prevent a connector from being run in the background. The -b option is *not* valid on the command-line invocation of connector_manager_*connector*.
-f*pollFrequency*	Poll frequency is the number of milliseconds between polling actions. • To specify the number of milliseconds, provide a value for *pollFrequency*. • To cause the connector to poll only when you type the value p in the connector's Command window, specify the -fkey option. • If a connector is configured to processes only business object requests and not application events, polling is unnecessary; you can disable polling by specifying -fno. The value of this parameter overrides any repository definitions. You can specify either -fkey or -fno, but not both. The -f option is valid on the command-line invocation of connector_manager_*connector*. The connector manager script can pass this option to its associated start_*connector*.sh script.

Table 12. Command-Line Options for the start_connector.sh Script (continued)

Option	Additional Information
`-t`*`threading_type`*	The *`threading_type`* parameter specifies the threading model: • `-tSINGLE_THREADED`: only a single thread accesses the application (The SAP connector uses `-tSINGLE_THREADED`.) • `-tMAIN_SINGLE_THREADED`: only the main thread accesses the application • `-tMULTI_THREADED`: multiple threads can access the application The `-t` option is *not* valid on the command-line invocation of `connector_manager_`*`connector`*. Specify it inside the generic `connector_manager` script, in the invocation of the `start_`*`connector`*`.sh` script.

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

You can also use the `connector_manager_`*`connector`* script to perform the following tasks:

- Stop a connector immediately by calling the API:

  ```
  connector_manager_connector -stop
  ```

- Show the current status of a connector:

  ```
  connector_manager_connector -stat
  ```

- Terminate the process by the process Id:

  ```
  connector_manager_connector -kill
  ```

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

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

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

```
connector_manager_connector.log
```

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

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

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

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

You can start the connector in several ways:

- click the desktop shortcut.

  Start the connector by double-clicking the desktop shortcut created as part of the installation procedure.
- Select the connector's menu option in the IBM WebSphere submenu of the Windows Start menu.
- Use a DOS Command Prompt window to execute the startup script.

  Open a DOS Command Prompt window and navigate to the appropriate connector directory. At the prompt, enter the one of the statements below, depending on whether the connector is a Java connector or a C++ connector:

  Java connector

  `start_Sap` *ConnectorName InterChangeServerName*

  C++ connector

  `start_connector` *ConnectorName InterChangeServerName*

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

Note: To figure out whether the connector is a Java connector or a C++ connector, navigate to

AdapterFrameworkProductDir`\documentation\ wbia_adapters\featurechecklists\versionlist.htm`

in your local directory, where *`AdapterFrameworkProductDir`* is the directory where you installed the WebSphere Business Integration Adapters product.

You can customize the startup for each connector by modifying the connector shortcut or the `start_`*`connector`*`.bat` file. Use the connector startup parameters listed in Table 13 to customize the startup of a connector.

Table 13. Connector Startup Parameters

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

Table 13. Connector Startup Parameters (continued)

Parameter	Function
`-f` *pollFrequency*	Poll frequency is the number of milliseconds between polling actions. • To specify the number of milliseconds, provide a value for *pollFrequency*. • To cause the connector to poll only when you type the value `p` in the connector's Command Prompt window, specify the `-fkey` option. • If a connector is configured to processes only business object requests and not application events, polling is unnecessary; you can disable polling by specifying `-fno`. The value of this parameter overrides any repository definitions. You can specify either `-fkey` or `-fno`, but not both.
`-j`	Specifies that the connector is written in Java. This parameter is optional if you specify `-l` *className.*
`-l` *className*	Specifies the name of the Java connector's global class, which is an extension of the connector base class. The startup script specifies this option for all Java connectors.
`-n` *connectorName*	Specifies the name of the connector to start.
`-p` password	Specifies the password that the connector uses to access InterChange Server.
`-s` *serverName*	Specifies the name of the InterChange Server. This parameter is required. The name is case-sensitive.
`-t`	Turns on the connector property `SingleThreadAppCalls`. This property guarantees that all calls the connector framework makes to the application-specific connector code are with one event-triggered flow. The default value is `false`. Important: Do not change the value of this property from its shipped value. Each connector has the appropriate setting for its threading model. Specify this option only when starting a connector you created.
`-x` *connectorProps*	Passes application-specific connector properties to the connector. Use the format `prop_name=`*`value`* for each value you enter.

Shutting down a connector

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

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

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

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

For instructions on enabling or disabling automatic restart, see "Steps for setting automatic and remote restart for a connector" on page 67.

Perform the following steps to shut down a connector:

UNIX

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

Windows

Use System Monitor to shut down a connector.

Instructions for shutting down a connector in Windows depends on which tool you use to monitor the system. The following sections detail steps for shutting down a connector using System Monitor or System Manager.

Steps for shutting down a connector in System Monitor: Perform the following steps in System Monitor to shut down a connector:

1. From the System Overview view, select the collaboration object of the connector you want to shut down by placing a check in the box to its left, then click the Pause icon from the upper-left corner of the view (see Figure 24 on page 59). Do this for each collaboration associated with the connector.
2. Select the connector you want to shut down by placing check in the box to its left, then click the Shutdown icon from the upper-left corner of the view (see Figure 24 on page 59).

Steps for shutting down a connector in System Manager: Perform the following steps in System Manager to shut down a connector:

1. From the expanded Collaboration Objects folder n the InterChange Server Component Management view of System Manager, (see Figure 27 on page 73), right-click the collaboration object associated with the connector, then choose Pause. Do this for each collaboration associated with the connector.

   The color on the collaboration object icon turns to yellow.
2. From the expanded Connectors folder in the InterChange Server Component Management view of System Manager, right-click the connector, then select Shut Down.

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

Restarting a connector

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

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

Steps for restarting a connector in System Monitor: Perform the following steps in System Monitor to restart a connector:

1. From the System Overview view (see Figure 23 on page 58), place a check in the box to the left of the connector you want to restart.
2. Click the Restart Agent icon from the upper-left corner of the view (see Figure 24 on page 59).

Steps for restarting a connector in System Manager: Perform the following steps in System Monitor to shut down a connector:

1. From the expanded Connectors folder in the InterChange System Component Management view of System Manager, (see Figure 29 on page 79), right-click the connector you want to restart.
2. Click Boot <`name_of_connector`>.

Steps for setting automatic and remote restart for a connector

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

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

Note: If the connector is already a member of a High Availability group, the automatic restart property would be redundant and should be disabled.

Perform the following steps to set up automatic and remote restart for a connector:

1. Install IBM WebSphere MQ.

   Use of the MQ-triggered OAD requires installation of the MQ-trigger Monitor and the configuration of certain queues. This monitor is installed as part of the WebSphere MQ software. These queues are created and configured by a special `mqtriggersetup.bat` script.

   Important: The WebSphere MQ-trigger Monitor must exist on the machine on which the connector agent resides. If multiple connector agents reside on a single machine, only one MQ-trigger Monitor needs to exist.
2. Start the MQ-trigger Monitor.
   To start the MQ-triggered OAD, you must start MQ-trigger Monitor, which can be done in either of the following ways:
   - Explicitly start MQ-trigger Monitor with the appropriate startup script.
   - Install MQ-trigger Monitor as a Windows service.
3. Configure a connector for the automatic and remote restart. Refer to "Steps for enabling connectors for MQ-triggered OAD" below.
4. Run the `mqtriggersetup.bat` script for each connector that needs to be restarted. (The `mqtriggersetup.bat` script is located in the `bin` directory.)

Steps for enabling connectors for MQ-triggered OAD: Perform the following steps to start Connector Configurator Express for the connector before you set the OAD properties:

1. Start InterChange Server Express.
2. Open System Manager.
3. Double-click the connector under Integration Component Libraries. This opens Connector Configurator Express.
4. On the Standard Properties tab, set the standard properties shown in Table 14.

Table 14. Configuring standard properties in Connector Configurator Express

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

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

- Initializing a connector for MQ-triggered OAD:

  Perform the following steps to enable automatic and remote restart for a connector for the first time:

  1. Set the `OADAutoRestartAgent` property to `True`.
  2. Set any of the other desired OAD properties in Table 14.
  3. Save the OAD properties in Connector Configurator Express.

- Toggling automatic and remote restart:

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

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

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

  – If you shut down the connector from within System Manager (by clicking the **Shutdown Agent** button), the connector agent shuts down. However, the

MQ-triggered OAD is not able to automatically restart the connector. You must restart the agent from within System Manager (by clicking the Reboot Agent button).

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

Steps for configuring flow control for connectors

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control system-wide, see“Steps for configuring system-wide flow control” on page 55. This section describes how to configure flow control for connectors.

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

To monitor how flow control is working in the system, you can view the Flow Control monitor and view provided as part of System Monitor or you can view the Statistics for collaboration objects or connectors from the InterChange Server Component Management view of System Manager. For more information on using the Flow Control monitor and view in System Monitor, see “Steps for reviewing default monitors” on page 2 and “Steps for using default views” on page 15. For more information on viewing the flow control from the InterChange Server Component Management view of System Manager, see “Collaboration object statistics” on page 26 or “Connector statistics” on page 27.

Perform the following steps to configure flow control for a connector, do the following:

1. In System Manger, navigate to the connector for which you want to configure flow control, then double-click that connector. Connector Configurator opens (see Figure 25 on page 70).

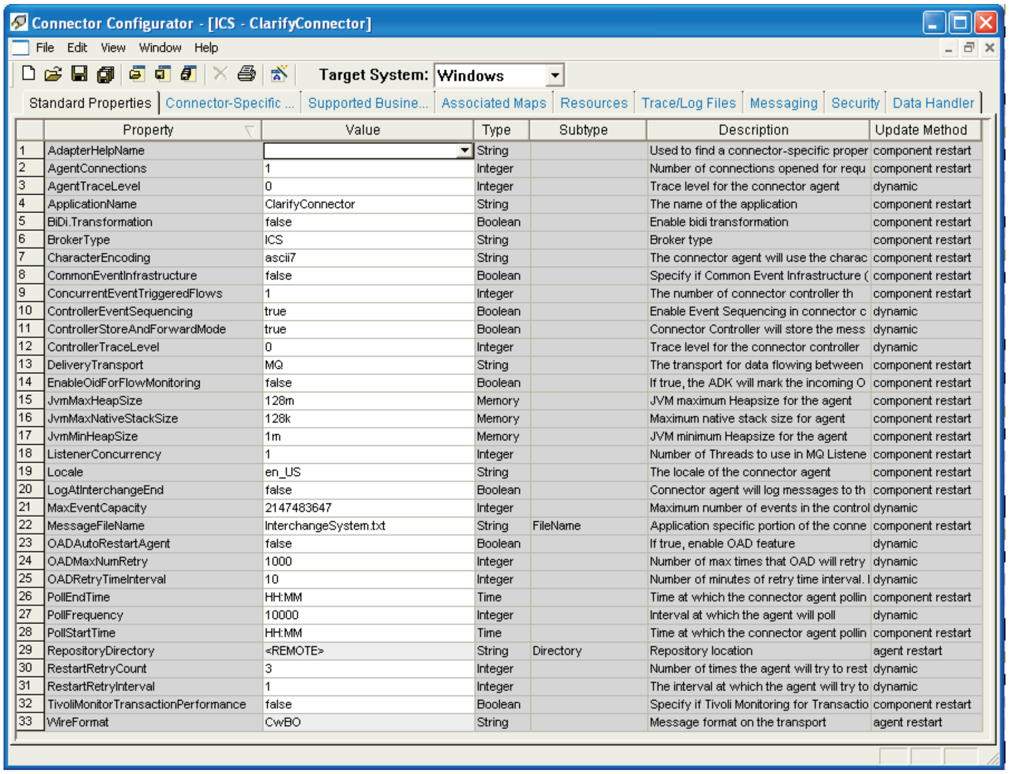

	Property	Value	Type	Subtype	Description	Update Method
1	AdapterHelpName		String		Used to find a connector-specific proper	component restart
2	AgentConnections	1	Integer		Number of connections opened for requ	component restart
3	AgentTraceLevel	0	Integer		Trace level for the connector agent	dynamic
4	ApplicationName	ClarifyConnector	String		The name of the application	component restart
5	BiDi.Transformation	false	Boolean		Enable bidi transformation	component restart
6	BrokerType	ICS	String		Broker type	component restart
7	CharacterEncoding	ascii7	String		The connector agent will use the charac	component restart
8	CommonEventInfrastructure	false	Boolean		Specify if Common Event Infrastructure (	component restart
9	ConcurrentEventTriggeredFlows	1	Integer		The number of connector controller th	component restart
10	ControllerEventSequencing	true	Boolean		Enable Event Sequencing in connector c	dynamic
11	ControllerStoreAndForwardMode	true	Boolean		Connector Controller will store the mess	dynamic
12	ControllerTraceLevel	0	Integer		Trace level for the connector controller	dynamic
13	DeliveryTransport	MQ	String		The transport for data flowing between	component restart
14	EnableOidForFlowMonitoring	false	Boolean		If true, the ADK will mark the incoming O	component restart
15	JvmMaxHeapSize	128m	Memory		JVM maximum Heapsize for the agent	component restart
16	JvmMaxNativeStackSize	128k	Memory		Maximum native stack size for agent	component restart
17	JvmMinHeapSize	1m	Memory		JVM minimum Heapsize for the agent	component restart
18	ListenerConcurrency	1	Integer		Number of Threads to use in MQ Listene	component restart
19	Locale	en_US	String		The locale of the connector agent	component restart
20	LogAtInterchangeEnd	false	Boolean		Connector agent will log messages to th	component restart
21	MaxEventCapacity	2147483647	Integer		Maximum number of events in the control	dynamic
22	MessageFileName	InterchangeSystem.txt	String	FileName	Application specific portion of the conne	component restart
23	OADAutoRestartAgent	false	Boolean		If true, enable OAD feature	dynamic
24	OADMaxNumRetry	1000	Integer		Number of max times that OAD will retry	dynamic
25	OADRetryTimeInterval	10	Integer		Number of minutes of retry time interval. I	dynamic
26	PollEndTime	HH:MM	Time		Time at which the connector agent pollin	component restart
27	PollFrequency	10000	Integer		Interval at which the agent will poll	dynamic
28	PollStartTime	HH:MM	Time		Time at which the connector agent pollin	component restart
29	RepositoryDirectory	<REMOTE>	String	Directory	Repository location	agent restart
30	RestartRetryCount	3	Integer		Number of times the agent will try to rest	dynamic
31	RestartRetryInterval	1	Integer		The interval at which the agent will try to	dynamic
32	TivoliMonitorTransactionPerformance	false	Boolean		Specify if Tivoli Monitoring for Transactio	component restart
33	WireFormat	CwBO	String		Message format on the transport	agent restart

Figure 25. Connector Configurator, Standard Properties tab

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

Administering JMS transport optimization

The flow of business information from adapters to the server, as well, as from the server to adapters, is a vital component of the WebSphere InterChange Server functionality. With the marked increase in the use of JMS Transport, enhancements were necessary to ensure the highest quality in performance, throughputs and scalability.

InterChange Server stores events in a persistent storage for recovery purposes. In a non-optimized state, this storage could be very costly, especially if the business object is expansive. In an optimized state, the event is left in the message queue and referenced in the database. When all event subscribers have completed their work, the message is deleted from the queue.

By synchronizing information in the critical sections, events can be retrieved sequentially from the queue, ensuring a retainable event sequence as well as a scalable server in a multi-processor environment.

In order to achieve JMS Transport optimization, the InterChange Server provides the following enhancements:

- **Improve caching** - queue objects are cached within the sender, enhancing adapter performance
- **Batch database operations** - business object events are accumulated in an ordered list and then persisted together in a batch operation, reducing the performance issues raised by frequent database operations
- **Optimize JMS recovery** - increase the event persistence performance, speed recovery operations and adapter response

This section covers the following topics:

"Optimization versus non-optimization" on page 71

"Steps for activating and de-activating optimization" on page 71

Optimization versus non-optimization

Although message transport is now optimized, there is also a need for transport to run in a non-optimized state , dependant upon business need. Switching from optimized to non-optimized allows users to swap messaging providers, if necessary, to accomodate the needs of their vendors.

You may opt to use a non-optimized state when business object events are small in size, or when database overhead is insignificant. However, prior to switching between an optimized and a non-optimized state you must wait until all queued events are recovered. Events running in an optimized state cannot be redelivered to InterChange Server in a non-optimized state.

Note: Optimization is designed to have minimal impact on in-bound service calls and Long Lived Business Process (LLBP) events, which will both continue to process as non-optimized events. This is possible since the optimized state can process both optimized and non-optimized events.

Steps for activating and de-activating optimization

Perform the following steps to activate and de-activate JMS transport optimization:

1. During the connector configuration, select the check box for JMS Optimization.
2. Set the value of the following connector properties. Once set, the connector configuration will upgrade the configuration files.
   - `jms.TransportOptimized` – `True` delivers events through optimized WIP.
   - `jms.ListenerConcurrency` – specifies the number of concurrent listeners for JMS transport. This property appears when`jms.TransportOptimized` is set to `True`.

   Note: If JMS is set as the transport, the default value for the `jms.TransportOptimized` property is `False`. When `jms.TransportOptimized` is set to `True`, the JMS provider (`jms.FactoryClassName`) must be IBM MQ.
3. To switch back to the non-optimized state, first ensure that the server is not currently processing any events and the delivery queue is clean. If you attempt to switch from an optimized state to a non-optimized state there are remaining events in the delivery queue, an error will display when the connector deploys to InterChange Server.
4. Clear the check box for JMS Optimization.

5. Set the value of the following connector properties. Once set, the connector configuration will upgrade the configuration files.
   - `jms.TransportOptimized` – `False` delivers events through non-optimized WIP.

Administering collaboration objects

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

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

This section covers the following topics:

“Viewing collaboration object states” on page 72

“Starting, stopping, and pausing collaboration objects” on page 74

“Configuring collaboration object run-time properties” on page 75

Viewing collaboration object states

You can view the state of a collaboration object either by logging on to System Monitor and opening a view that contains collaboration object information or by using the InterChange Server Component Management view of System Manager. To log on to System Monitor, follow the instructions in “Steps for logging on to System Monitor” on page 12.. To use the InterChange Server Component Management view of System Manager, follow the instructions in “Steps for connecting to an InterChange Server instance” on page 24..

The state of a collaboration object is represented differently, depending on whether you ar e using System Monitor or System Manager.

Steps for using System Monitor to view collaboration objects

Perform the following steps to view the state of collaboration objects using System Monitor:

1. If the System Overview view is not displayed, click the System Overview link under Views in the left pane of the Web page. The System Overview Monitor appears (see Figure 11 on page 13) in the body of the Web page.

   When the product is installed, the default view is set to System Overview, and the default monitor contained in that view is set to System Overview. These defaults can be changed to suit your monitoring needs. See “Setting up views to monitor the system” on page 14 for instructions.
2. Click the triangle next to the name of the server to reveal a list of components on the system. All collaboration objects are listed along with their status, start time, and total up time (see Figure 26 on page 73).

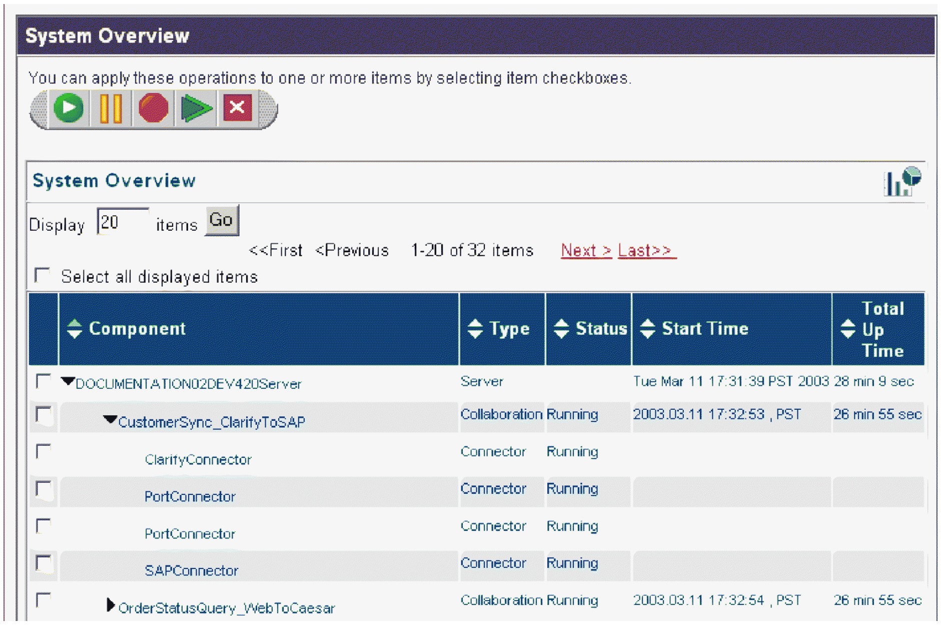

System Overview

You can apply these operations to one or more items by selecting item checkboxes.

System Overview

Display 20 items Go

<<First <Previous 1-20 of 32 items Next > Last>>

□ Select all displayed items

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

Figure 26. System Monitor, System Overview displaying collaboration object status

Note: You may also view collaboration states using the Collaboration Overview view.

Steps for using System Manager to view collaboration object states

Perform the following steps to view the state of collaboration objects using System Manager:

1. Connect to the InterChange Server instance that contains the collaboration object you want to view. See "Steps for connecting to an InterChange Server instance" on page 24 for instructions on connecting to an InterChange Server instance.
2. Expand the InterChange Server instance, then expand the Collaboration Objects folder.

   The collaboration objects appear under the expanded Collaboration Objects folder with different colored lights to indicate their different states (see Figure 27.

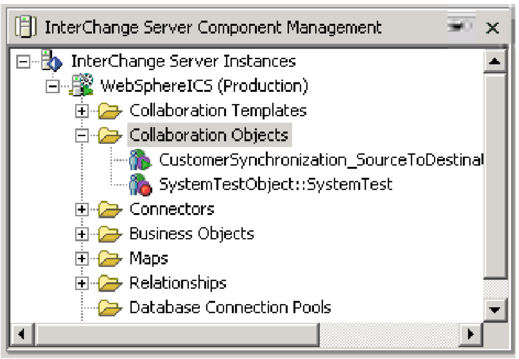

InterChange Server Component Management

- InterChange Server Instances
  - WebSphereICS (Production)
    - Collaboration Templates
    - Collaboration Objects
      - CustomerSynchronization_SourceToDestinat
      - SystemTestObject::SystemTest
    - Connectors
    - Business Objects
    - Maps
    - Relationships
    - Database Connection Pools

Figure 27. Collaboration Objects folder in the InterChange Server Component Management view of System Manager

Table 15 describes the collaboration object states which are viewable from System Monitor and the Collaboration Objects folder in the InterChange Server Component Management view of System Manager:

Table 15. Collaboration object states

Collaboration object state	Description
Start	Starting a collaboration object causes it to subscribe to its triggering business objects and to process them as they arrive. If you stop and then restart InterChange Server, collaboration objects in the Start state automatically start running when InterChange Server comes back up.
Pause	Pausing a collaboration prevents it from receiving new flow initiators. The collaboration completes all of the current processing, then enters an idle state. A connector maintains its subscription information; therefore, it continues to send flow initiators to the connector queues. The collaboration processes these when it is resumed. To resume collaboration execution, click Start in System Monitor or in the Collaboration Object menu of the InterChange Server Component Management view of System Manager.
Stop	Stopping a collaboration causes it to unsubscribe to business objects. The collaboration completes all of the current processing, then becomes inactive. Unlike the Pause command, the Stop command causes connectors to stop sending business objects to the collaboration. To properly stop a collaboration without losing any flows, first stop the associated connectors from polling, allow all flows to process, then stop the collaboration.
Shut Down	Shutting down a collaboration immediately ends processing of current flows. When the collaboration is restarted, the system recovers by processing those flows that were interrupted by the shut down and recovering those flows waiting in the queue. This recovery is not immediate, so prepare to wait while the system completes the recovery interval.

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

Starting, stopping, and pausing collaboration objects

To make a collaboration object functional for the first time, you must first configure it then start it. See “Configuring collaboration object run-time properties” on page 75 for more information on configuring collaborations. Depending on which tool you are using, you run, stop, and pause collaboration objects in different ways.

Steps for using System Monitor to start, stop and pause collaboration objects

Perform the following steps to use System Monitor to start, stop and pause collaboration objects:

1. While viewing the System Overview view (see Figure 23 on page 58), select a collaboration object by placing check in the box to its left.
2. Select the Start, Pause, or Stop icon from the icon group in the upper-left corner of the view (see Figure 24 on page 59).

Note: You may also start, stop, pause and shutdown collaboration objects using the Collaboration Overview view.

Steps for using System Manager to start, stop and pause collaboration objects

Perform the following steps to use System Manager to start, stop and pause collaboration objects:

1. From the Collaboration Objects folder of the InterChange Server Component Management view, right-click a collaboration object.
2. Select the Start, Pause, or Stop.

Configuring collaboration object run-time properties

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

"Steps for setting collaboration object general properties" on page 75

"Steps for configuring collaboration objects to process concurrent event-triggered flows" on page 77

"Steps for configuring flow control for collaboration objects" on page 77

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

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

Steps for setting collaboration object general properties

Perform the following steps to open the Collaboration Properties window and change values for general properties of a collaboration object, do this:

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

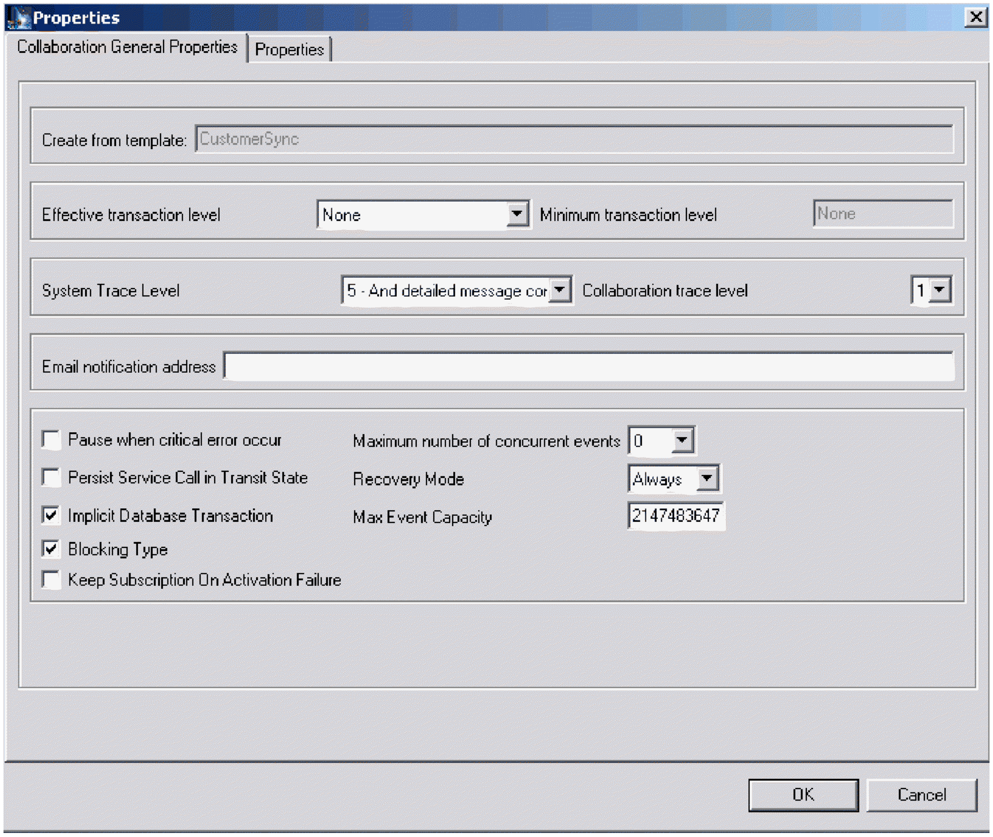

Create from template: CustomerSync

Effective transaction level: None

Minimum transaction level: None

System Trace Level: 5 - And detailed message cor

Collaboration trace level: 1

Email notification address:

☐ Pause when critical error occur

☐ Persist Service Call in Transit State

☑ Implicit Database Transaction

☑ Blocking Type

☐ Keep Subscription On Activation Failure

Maximum number of concurrent events: 0

Recovery Mode: Always

Max Event Capacity: 2147483647

Figure 28. Properties dialog box, Collaboration General Properties tab

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

The dialog box enables you to make settings for the following:

- Effective Transaction Level

  See “Setting the Effective Transaction Level and Other General Properties” in the *System Implementation Guide*.
- System Trace Level

  See the *IBM WebSphere InterChange Server Problem Determination Guide*.
- Collaboration Trace Level

  See the *IBM WebSphere InterChange Server Problem Determination Guide*.
- E-mail Notification Address

  See See the *IBM WebSphere InterChange Server Problem Determination Guide*.
- Pause When Critical Error Occurs

  See the *IBM WebSphere InterChange Server Problem Determination Guide*..
- Persist Service Call In Transit State

  See Chapter 3, “Administering problem scenarios,” on page 131, and see the *Collaboration Development Guide*.
- Implicit Database Transaction

  See “Configuring Transaction Bracketing” in the *System Implementation Guide*.
- Blocking Type

  See "Blocking type" in the *System Implementation Guide*.

- Maximum Number of Concurrent Events

  See "Steps for configuring collaboration objects to process concurrent event-triggered flows" on page 77.
- Recovery Mode

  See the *IBM WebSphere InterChange Server Problem Determination Guide*.
- Max Event Capacity

  See the *IBM WebSphere InterChange Server Problem Determination Guide*.

Steps for configuring collaboration objects to process concurrent event-triggered flows

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

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

Perform the following steps to set the maximum number of concurrent flows for a collaboration:

1. From the expanded Collaboration Objects folder in the InterChange Server Component Management view, right-click the collaboration object that you want to change, then select Properties. The Properties dialog box appears (see Figure 28 on page 76).
2. In the Collaboration General Properties tab, enter a value in the ″Maximum number of concurrent events″ field.
3. Click OK to save your changes and close the window.
4. Restart the collaboration for changes to take effect.

Steps for configuring flow control for collaboration objects

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control system-wide, see"Steps for configuring system-wide flow control" on page 55. This section describes how to configure flow control for collaboration objects.

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

To monitor how flow control is working in the system, you can view the Flow Control monitor and view provided as part of System Monitor or you can view the Statistics for collaboration objects or connectors from the InterChange Server Component Management view of System Manager. For more information on using the Flow Control monitor and view in System Monitor, see "Steps for reviewing default monitors" on page 2 and "Steps for using default views" on page 15. For more information on viewing the flow control from the InterChange Server Component Management view of System Manager, see "Collaboration object statistics" on page 26 or "Connector statistics" on page 27.

Perform the following steps to configure flow control for a collaboration object:

1. From the expanded Collaboration Objects folder in the InterChange Server Component Management view of System Manager, right-click the collaboration object for which you want to create flow control, then select Properties from the drop-down menu. The Properties dialog box appears (see Figure 28 on page 76).
2. In the Collaboration General Properties tab, edit the value in the Max Event Capacity field to represent the maximum number of events you want queued for a collaboration object. The valid range of values for this property is from 1 to 2147483647.
3. Click OK. The property is changed immediately.

Steps for reconfiguring the timeout attribute for long-lived business processing

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

Perform the following steps to reconfigure the service call timeout values of a collaboration with long-lived business processing:

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

   Note: The name of the service call timeout configuration property may be something like, ″CreateTimeout″ or ″RetreiveTimeout,″ but since there is no naming convention for this property, you may have to contact the person who developed the collaboration, if the name of the service call timeout configuration property is not immediately apparent.
3. Edit the value so that it represent the number of timeout minutes allowed.

   Note: The Value field must contain an integer greater than 0. If it contains a 0 or is left blank, the waittime is equal to infinity. If it contains non-numerical values, it will trigger a collaboration run time exception.
4. Click OK. Your changes take place immediately, without the need to restart InterChange Server.

Administering maps

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

This section covers the following topics:

“Viewing map states” on page 79

“Starting and stopping maps” on page 80

Viewing map states

You can view the state of a map either by logging on to System Monitor and opening a view that contains map status or by using the InterChange Server Component Management view of System Manager. To log on to System Monitor, follow the instructions in "Steps for logging on to System Monitor" on page 12. To use the InterChange Server Component Management view of System Manager, follow the instructions in "Steps for connecting to an InterChange Server instance" on page 24..

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

Steps for viewing map states in System Monitor

Perform the following steps to view the state of map using System Monitor:

1. With the System Overview view displayed, click the Maps and Relationships link under Views in the left pane of the Web page. The Map Status and Relationship Status monitors appear (see Figure 11 on page 13) in the body of the Web page.
2. The default view and default monitor are set to System Overview. These defaults can be changed to suit your monitoring needs. See "Setting up views to monitor the system" on page 14 for instructions.

Note: You may also view map states using the Maps and Relationships view.

Steps for viewing map states in System Manager

Perform the following steps to view the state of map using System Manager:

1. Connect to the InterChange Server instance that contains the map you want to view. See "Steps for connecting to an InterChange Server instance" on page 24 for instructions on connecting to an InterChange Server instance.
2. Expand the InterChange Server instance, then expand the Maps folder. The maps appear under the expanded Maps folder with different colored lights to indicate their different states (see Figure 29).

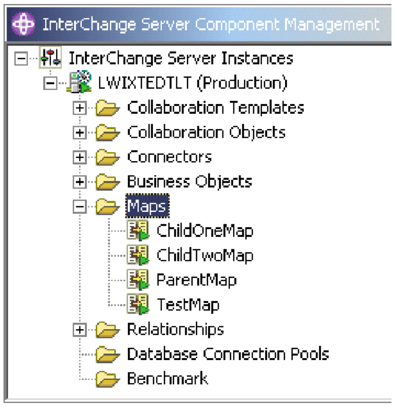

Figure 29. Maps folder in the InterChange Server Component Management view of System Manager

Table 16 lists the map states represented by the display color on the map icon and describes what actions can be performed during that state. For more information about maps, refer to the *Map Development Guide.*

Table 16. Map States

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

Starting and stopping maps

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

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

Map activation

For a map to be executable, it must first be activated. Map Designer automatically starts a map when it successfully compiles the map. However, other changes to the map might require that you explicitly stop and restart the map for the change to take effect. See Appendix B, “Requirements for restarting IBM WebSphere Business Integration system components,” on page 157 to find out what changes require maps to be stopped and restarted.

Steps for starting and stopping maps

Depending on which tool you are using, you start and stop maps in different ways.

Steps for starting and stopping maps in System Monitor

Perform the following steps to start or stop a map using System Monitor:

1. If the System Overview view is displayed, click the Maps and Relationships link under Views in the left pane of the Web page. The Map Status and Relationship Status monitors appear (see Figure 11 on page 13) in the body of the Web page.
2. Select the Start or Stop icon from the icon group in the upper-left corner of the view (see Figure 24 on page 59).

Note: You may also start and stop map using the Maps and Relationships view.

Steps for starting and stopping maps in System Manager

Perform the following steps to start or stop a map using System Manager:

1. From the expanded Maps folder in the InterChange Server Component Management view (see Figure 29 on page 79), right.-click a map.
2. Select either the Start <name_of_map> or Stop <name_of_map> option.

Administering relationships

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

This section covers the following topics:

"Viewing relationship states" on page 81

"Starting and stopping relationships" on page 82

Viewing relationship states

You can view the state of a relationship either by logging on to System Monitor and opening a view that contains relationship status or by using the InterChange Server Component Management view of System Manager. To log on to System Monitor, follow the instructions in "Steps for logging on to System Monitor" on page 12. To use the InterChange Server Component Management view of System Manager, follow the instructions in "Steps for connecting to an InterChange Server instance" on page 24..

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

Steps for viewing relationship states in System Monitor

Perform the following steps to view the state of relationships using System Monitor:

1. With the System Overview view displayed, click the Maps and Relationships link under Views in the left pane of the Web page. The Map Status and Relationship Status monitors appear (see Figure 11 on page 13) in the body of the Web page.

The default view and monitor are set to System Overview. These defaults can be changed to suit your monitoring needs. See "Setting up views to monitor the system" on page 14 for instructions.

Steps for viewing relationship states in System Manager

Perform the following steps to view the state of relationships using System Manager:

1. Connect to the InterChange Server instance that contains the relationship you want to view. See "Steps for connecting to an InterChange Server instance" on page 24 for instructions on connecting to an InterChange Server instance.
2. Expand the InterChange Server instance, then expand the Relationships folder, then expand either the Dynamic or Static folder. The relationships appear under either the expanded Dynamic folder or the expanded Static folder, and have different colored lights to indicate their different states (see Figure 30 on page 82).

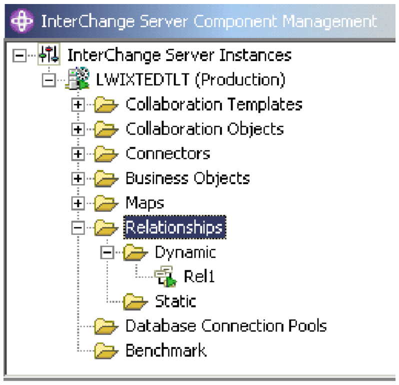

Figure 30. Relationships folder in the InterChange Server Component Management view of System Manager

Table 17 lists the relationship states represented by the display color and describes what actions can be performed during that state.

Table 17. Relationship States

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

Starting and stopping relationships

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

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

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

This section describes the following topics:

"Relationship activation" on page 83

"Steps for starting and stopping relationships"

"Relationship table caching"

Relationship activation

For a relationship to be executable, it must be activated. However, you cannot modify a relationship when it is active. Therefore, you must stop the relationship, make the change to the relationship, and then restart the relationship. See Appendix B, "Requirements for restarting IBM WebSphere Business Integration system components," on page 157 to find out what changes require relationships to be restarted.

Steps for starting and stopping relationships

Depending on which tool you are using, you start and stop relationships in different ways.

Steps for starting and stopping relationships using System Monitor

Perform the following steps to start and stop relationships using System Monitor:

1. From the System View window (see Figure 11 on page 13), select Maps and Relationships from the View drop-down menu. The Map Status and Relationship Status monitors appear.
2. Click the checkbox for the relationship you want to start or stop.
3. Click the **Start** or **Stop** button to perform the appropriate action.

Steps for starting and stopping relationships using System Manager

Perform the following steps to start and stop relationships using System Manager:

From the expanded Dynamic or Static folder in the InterChange Server Component Management view of System Manager (see Figure 30 on page 82), right-click the name of a relationship, then select Start `<relationship_name>` or Stop `<relationship_name>`.

Relationship table caching

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

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

Steps for enabling caching: Perform the following steps to enable relationship table caching for a static relationship:

1. Expand the Relationships folder in System Manager.
2. Expand the Static folder in the object browser to locate the static relationship whose tables you want to be cached.

3. Right-click the static relationship to determine its current cached state. If the Cached option appears with no check mark to the left, caching for that relationship is currently disabled. Choose Cached from the context menu to enable caching.

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

Steps for disabling caching: Perform the following steps to disable relationship table caching for a static relationship:

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

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

Steps for reloading the cached tables: Perform the following steps to have InterChange Server reread the relationship's tables into memory with the Reload feature:

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

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

Steps for tracing cached tables: Perform the following step to have InterChange Server log a trace message each time it loads and unloads relationship tables in memory:

1. Set the `RELATIONSHIP.CACHING` configuration parameter to five (5) in the `TRACING` section of the `InterchangeSystem.cfg` file:

   ```
   RELATIONSHIP.CACHING=5
   ```

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

Using Relationship Manager

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

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

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

Figure 31 shows a sample of a relationship tree in Relationship Manager for an identity relationship.

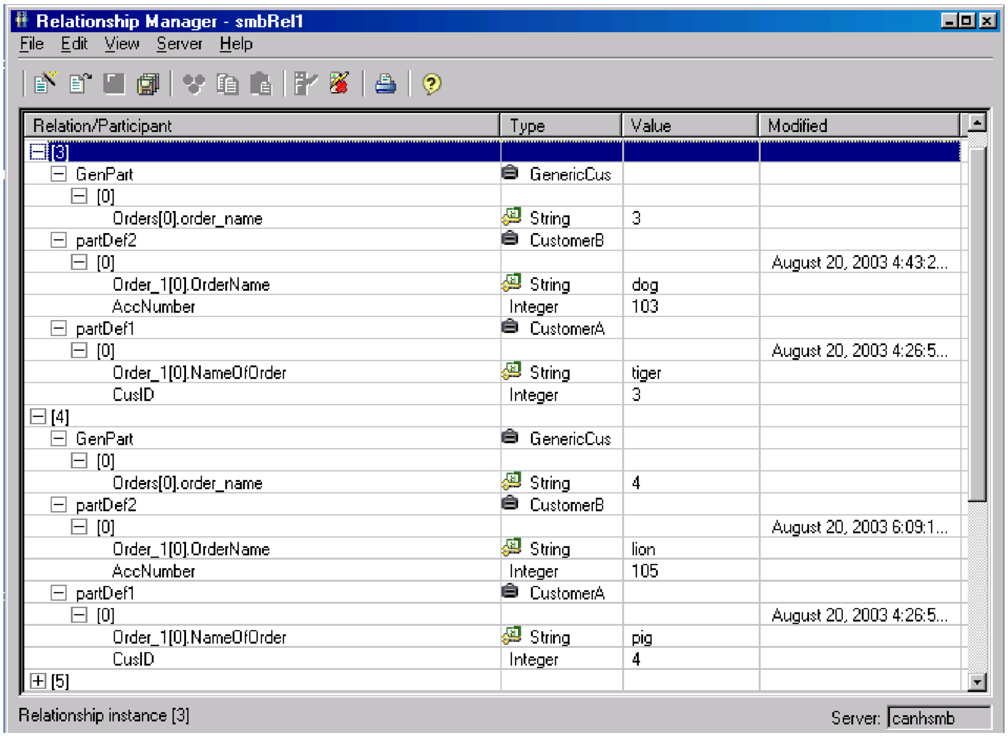

Relation/Participant	Type	Value	Modified
⊟ [3]			
⊟ GenPart	GenericCus		
⊟ [0]			
Orders[0].order_name	String	3	
⊟ partDef2	CustomerB		
⊟ [0]			August 20, 2003 4:43:2...
Order_1[0].OrderName	String	dog	
AccNumber	Integer	103	
⊟ partDef1	CustomerA		
⊟ [0]			August 20, 2003 4:26:5...
Order_1[0].NameOfOrder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

Relationship instance [3]

Server: canhsmb

Figure 31. Relationship Manager, relationship tree

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

- Create and delete relationship instances

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

This section covers the following topics:

"Steps for starting Relationship Manager"

"Connecting to and disconnecting from a server" on page 87

"Connecting to and disconnecting from a server" on page 87

"Working with relationships in Relationship Manager" on page 88

"Working with relationship data" on page 95

Steps for starting Relationship Manager

Perform one of the following steps to start Relationship Manager:

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

Relationship Manager starts. At this point it is disconnected from the server; you must connect to an InterChange Server instance as described in "Steps for connecting to Relationship Manager from InterChange Server" on page 87 to proceed further.

Figure 32 on page 87 shows Relationship Manager in a disconnected state.

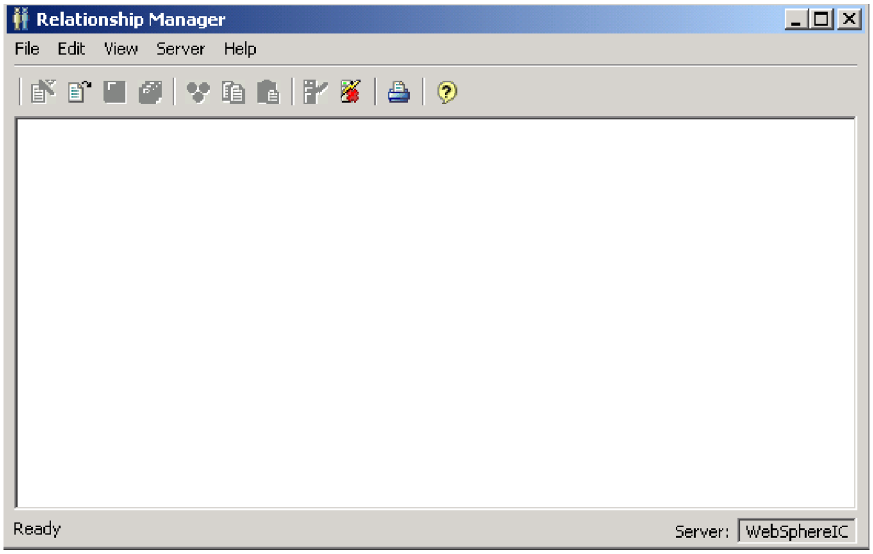

Server: WebSphereIC

Figure 32. Relationship Manager

Connecting to and disconnecting from a server

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

- "Steps for connecting to Relationship Manager from InterChange Server"
- "Steps for disconnecting from InterChange Server" on page 88

Steps for connecting to Relationship Manager from InterChange Server

Perform the following steps to connect Relationship Manager to InterChange Server:

1. Select **Server > Connect** from the menu bar of Relationship Manager. The Connect to InterChange Server dialog box appears (see Figure 33).

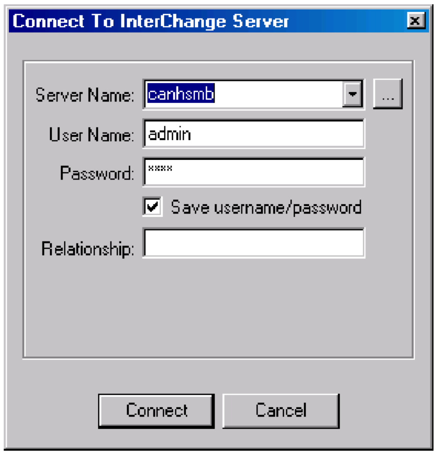

Server Name: canhsmb

User Name: admin

Password: ****

☑ Save username/password

Relationship:

Figure 33. Connect to InterChange Server

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

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

   If you do not want to open a relationship at this time you can open it after connecting to the server. For more information, see "Steps for opening a relationship" on page 89.
7. Click **Connect**.

   If you connect to InterChange Server in Relationship Manager and specify a relationship to open as described in step 6, then Relationship Manager opens the Retrieve Relationship Instances window, described in "Steps for retrieving relationship instances" on page 89.

Steps for disconnecting from InterChange Server

Perform the following steps to disconnect Relationship Manager from InterChange Server:

1. Select **Server > Disconnect** from the menu bar of Relationship Manager.

Working with relationships in Relationship Manager

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

"Steps for opening a relationship" on page 89

"Steps for retrieving relationship instances" on page 89

"Steps for creating relationship instances" on page 92

"Steps for deleting relationship instances" on page 93

"Deactivating and activating participants" on page 93

"Steps for copying participants" on page 94

"Steps for loading and unloading business object files" on page 94

Steps for opening a relationship

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

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

   Figure 34 shows the "Open Relationship" window.

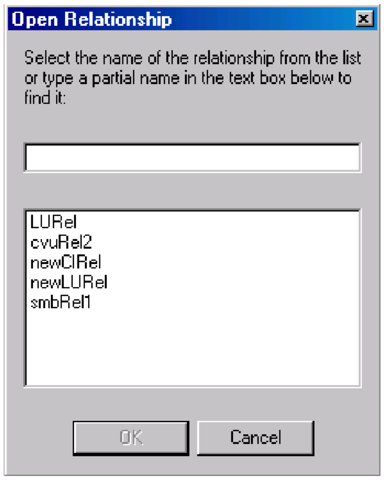

Open Relationship

Select the name of the relationship from the list or type a partial name in the text box below to find it:

LURel
cvuRel2
newClRel
newLURel
smbRel1

Figure 34. Opening a relationship

3. Click **OK**.

   When you open a relationship, Relationship Manager opens the Retrieve Relationship Instances window, described in "Steps for retrieving relationship instances."

Steps for retrieving relationship instances

Perform the following step to retrieve relationship instances:

1. Select **File > Retrieve** from the menu bar of Relationship Manager to retrieve relationship instances or return a count of how many instances there are for a relationship. Figure 35 on page 90 shows the Retrieve Relationship Instances window.

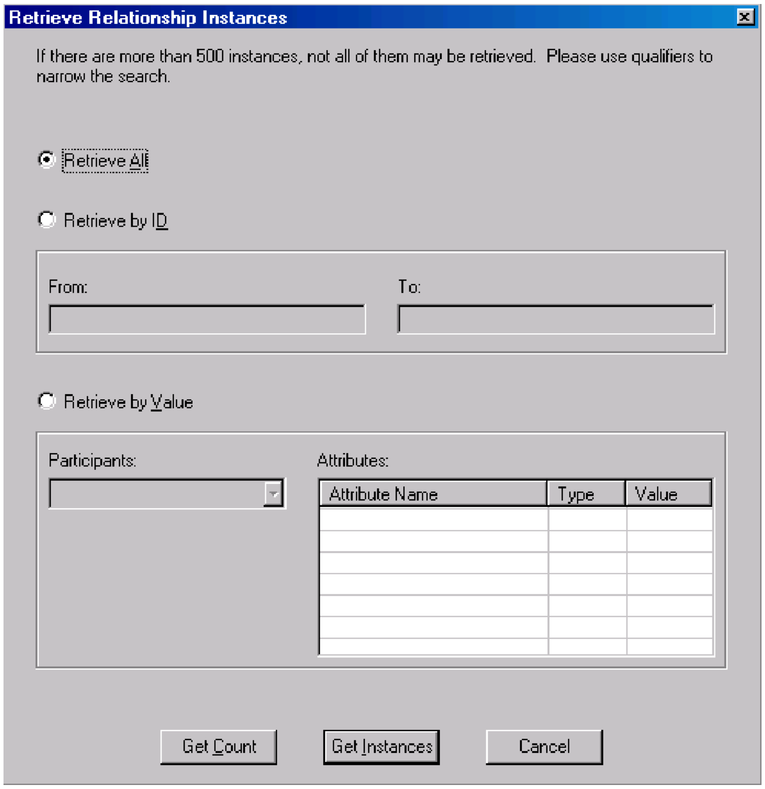

Retrieve Relationship Instances

If there are more than 500 instances, not all of them may be retrieved. Please use qualifiers to narrow the search.

● Retrieve All

○ Retrieve by ID

From:

To:

○ Retrieve by Value

Participants:

Attributes:

Attribute Name	Type	Value

Figure 35. Retrieving relationship instances

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

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

- Retrieve the first 500 instances for the relationship, as described in “Steps for retrieving all instances.”
- Retrieve a range of instances for the relationship based on the relationship instance ids, as described in “Steps for retrieving by relationship ID.”
- Retrieve a relationship instance that contains a participant of a particular value you specify, as described in “Steps for retrieving by participant data” on page 91.
- Return a count of the number of instances for the relationship, as described in “Steps for returning a count of the relationship instances” on page 91.

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

Steps for retrieving all instances: Perform the following steps to retrieve the first 500 instances for a relationship:

1. At the Retrieve Relationship Instances window click **Retrieve All**.
2. Click **Get Instances**.

   Relationship Manager displays the first 500 instances for the relationship.

Steps for retrieving by relationship ID: Perform the following steps to retrieve a range of up to 500 instances by relationship ID:

1. At the Retrieve Relationship Instances window click **Retrieve by ID**.
2. Type the ID of the first instance in the range you want to retrieve in the **From** field.
3. Type the ID of the last instance in the range you want to retrieve in the **To** field.
4. Click **Get Instances**.

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

Steps for retrieving by participant data: Perform the following steps to retrieve a relationship instance based on values for key or non-key attributes of selected participants:

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

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

   For lookup relationships, the drop-down menu lists the participant names followed by the word "Data".
3. Type one of the types of values listed in Table 18 in the **Value** column in the "Attributes" pane.

Table 18. Supported values for retrieving relationship instances by participant data

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

4. Click **Get Instances**.

   Relationship Manager displays the first 500 relationship instances that match the specified value.

Steps for returning a count of the relationship instances: Perform the following steps to return the number of relationship instances that satisfy a retrieval criteria:

1. Select the options for the criteria as described in "Steps for retrieving all instances" on page 90, "Steps for retrieving by relationship ID" on page 90, or "Steps for retrieving by participant data"
2. Select **Get Count**.

Relationship Manager displays the first 500 instances for the relationship.

Steps for creating relationship instances

Perform the following steps to create a new instance for a relationship:

1. Create the new relationship instance by performing any of the following tasks:
   - Select **File > New Instances** from the menu bar.
   - Use the keyboard shortcut **Ctrl+N**.
   - Click the **New Relationship Instance** icon in the toolbar.

   Relationship Manager displays the new relationship instance.

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

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

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

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

Table 19. Tasks for Participant Data

Task	Action
Save the participant instance.	To save the new participant instance, right-click the participant instance and choose **Save Participant** from the context menu. Relationship Manager saves in the appropriate relationship table the data for this participant. The **Modified** column for the participant instances displays the date the participant was saved, which is the create date, in this case. **Note:** Once the participant data has been saved, it cannot be changed. To change its data, the participant must be deleted and another created.

Table 19. Tasks for Participant Data (continued)

Task	Action
Add more participant instances.	Repeat repeat step 3 on page 92 in the previous list. **Note:** If you are working with an identity relationship, you cannot create more than one participant instance for a participant definition.
Delete a participant.	If necessary, you can delete a saved participant instance by right-clicking on the participant instance and choosing **Delete Participant** from the context menu. Relationship Manager removes the participant instance from the relationship table. If you do not want to remove the participant instance from the database, use the **Deactivate Participant** option (see "Deactivating and activating participants" on page 93). A deactivated participant retains its instance ID and its values.
Save the relationship instance.	Save the relationship instance by performing one of the following tasks: • Select **File > Save** from the menu bar (activated when a relationship instance is selected). • Right-click the relationship instance and choose **Save Relationship** from the context menu. InterChange Server generates the relationship instance ID and Relationship Manager replaces the ??? placeholder with this new ID. Relationship Manager updates the modified date on all saved participant instances to this date. **Note:** At least one participant instance and all key attribute data must be created before the relationship instance can be saved.
Save all relationship instances.	Select **File > Save All** from the menu bar. InterChange Server generates the relationship instance IDs for any relationship instances that do not have one. Relationship Manager replaces any "???" placeholders with the new IDs. Relationship Manager updates the modified date on all saved participant instances to this date.

Steps for deleting relationship instances

Perform the following steps to delete a relationship instance from the relationship tables:

1. Select the relationship instance you want to delete.
2. Do one of the following in Relationship Manager:
   - Select **File > Delete** from the menu bar.
   - Right-click the relationship instance and select **Delete Relationship Instance** from the context menu.

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

Deactivating and activating participants

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

Steps for deactivating a participant: Perform the following steps to deactivate a participant instance:

1. Right-click the participant instance you want to deactivate.
2. Select **Deactivate Participant** from the context menu.

The participant is removed from the Relationship Manager display but not from the relationship tables.

Steps for activating a participant: Perform the following steps to activate a participant instance:

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

   The Deactivated Participants window displays as shown in Figure 36.

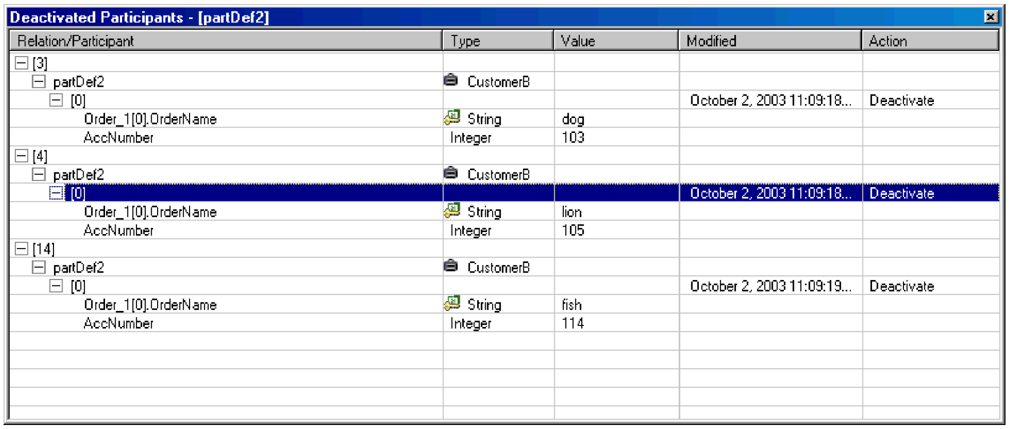

Relation/Participant	Type	Value	Modified	Action
[3]				
partDef2	CustomerB			
[0]			October 2, 2003 11:09:18...	Deactivate
Order_1[0].OrderName	String	dog		
AccNumber	Integer	103		
[4]				
partDef2	CustomerB			
[0]			October 2, 2003 11:09:18...	Deactivate
Order_1[0].OrderName	String	lion		
AccNumber	Integer	105		
[14]				
partDef2	CustomerB			
[0]			October 2, 2003 11:09:19...	Deactivate
Order_1[0].OrderName	String	fish		
AccNumber	Integer	114		

Figure 36. Deactivating participants

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

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

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

Steps for copying participants

You can create a new participant instance by copying an existing participant instance. Perform the following steps to copy a participant instance:

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

Steps for loading and unloading business object files

You can load a business object file of the same type into a participant. Perform the following steps to load a business object data file into a participant:

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

The Participant window displays the business object associated with that participant instance, as shown in Figure 37.

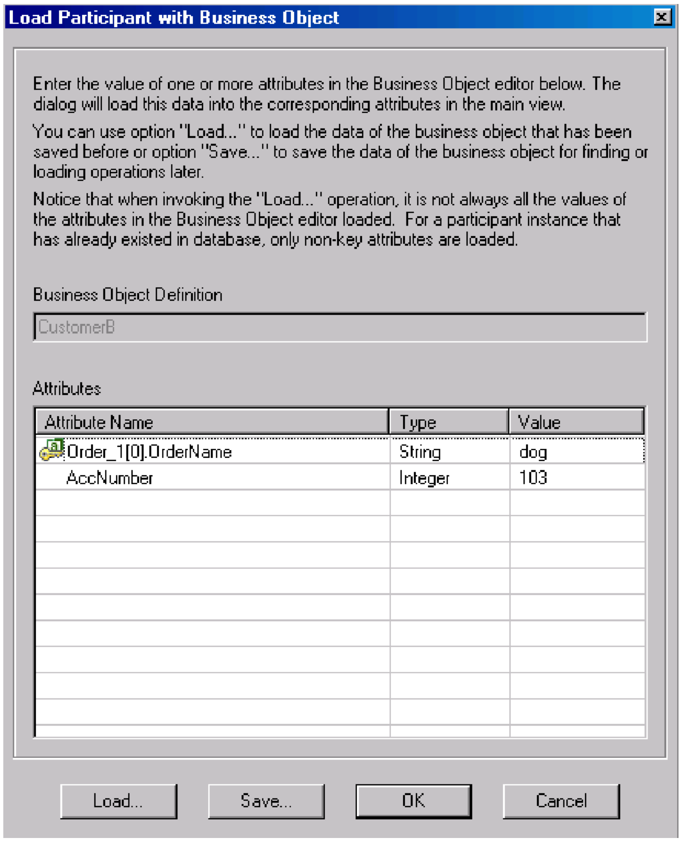

Enter the value of one or more attributes in the Business Object editor below. The dialog will load this data into the corresponding attributes in the main view.

You can use option "Load..." to load the data of the business object that has been saved before or option "Save..." to save the data of the business object for finding or loading operations later.

Notice that when invoking the "Load..." operation, it is not always all the values of the attributes in the Business Object editor loaded. For a participant instance that has already existed in database, only non-key attributes are loaded.

Business Object Definition

CustomerB

Attributes

Attribute Name	Type	Value
Order_1[0].OrderName	String	dog
AccNumber	Integer	103

Figure 37. Loading participants with business objects

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

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

Working with relationship data

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

“Steps for searching for participants” on page 96

“Steps for filtering the displayed participants” on page 97

“Steps for cleaning up participants” on page 98

“Steps for printing relationship data” on page 99

Steps for searching for participants

You can search for participant instances based on different criteria. Depending on how specific your search criteria is, your searches can locate a unique participant instance or a group of participant instances. You can find participant instances either by business object or by data.

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

Perfom the following steps to search for instances by business object:

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

   The Find Instances by Business Object window appears (see Figure 38).

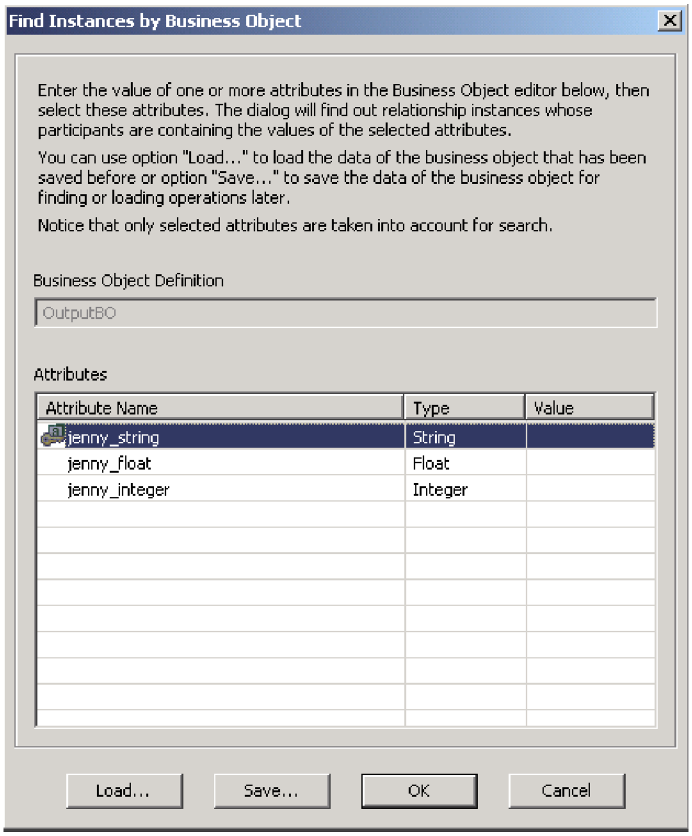

Find Instances by Business Object

Enter the value of one or more attributes in the Business Object editor below, then select these attributes. The dialog will find out relationship instances whose participants are containing the values of the selected attributes.

You can use option "Load..." to load the data of the business object that has been saved before or option "Save..." to save the data of the business object for finding or loading operations later.

Notice that only selected attributes are taken into account for search.

Business Object Definition: OutputBO

Attributes

Attribute Name	Type	Value
jenny_string	String	
jenny_float	Float	
jenny_integer	Integer	

Figure 38. Find Instances by Business Object

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

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

Steps for finding instances by data: This option searches for instances whose type is Data.

Perfom the following steps to search for instances by data:

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

   Relationship Manager displays the "Find Instances by Data" window, as shown in Figure 39.

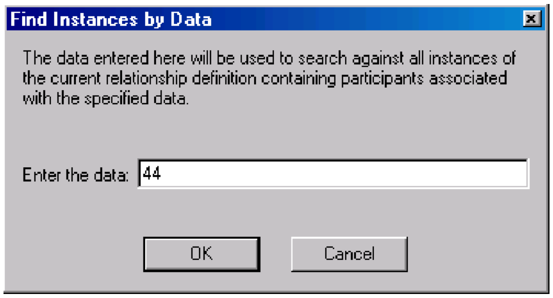

The data entered here will be used to search against all instances of the current relationship definition containing participants associated with the specified data.

Enter the data: 44

Figure 39. Finding instances by data

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

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

Steps for filtering the displayed participants

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

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

   Relationship Manager displays the "Filter" dialog box, as shown in Figure 40.

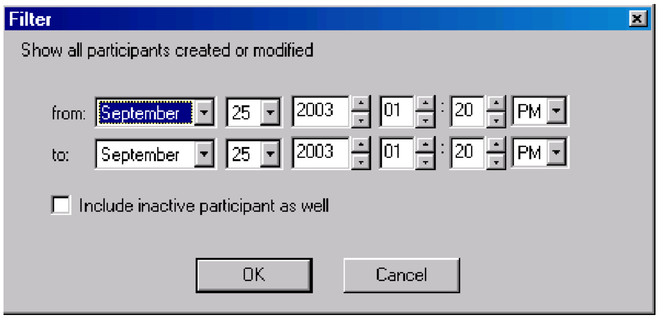

Show all participants created or modified

from: September 25 2003 01 : 20 PM

to: September 25 2003 01 : 20 PM

□ Include inactive participant as well

Figure 40. Filtering participant results

3. In the "Filter" dialog box, enter the earliest date of creation or modification for the participant in the **from** field and the latest date of creation or modification for the participant in the **to** field.

   Use the following techniques to enter the date value:
   - Type letters in the text field to cycle through days of the week. For instance, type **S** to cycle through dates that occur on Saturdays and Sundays.
   - Click the small up and down arrows to increment or decrement the date by one day.

- Click the large down arrow to display a calendar that you can use to select a date.

4. Select the **Include inactive participants as well** check box if you want to include inactive participants in the resulting display.
5. Click **OK**.

   Relationship Manager displays the history of activity for the filtered interval in the “Filter Results” dialog box. The filtered display includes inactive participants if the **Include inactive participants as well** option box is checked. Figure 41 shows the “Filter Results” dialog box.

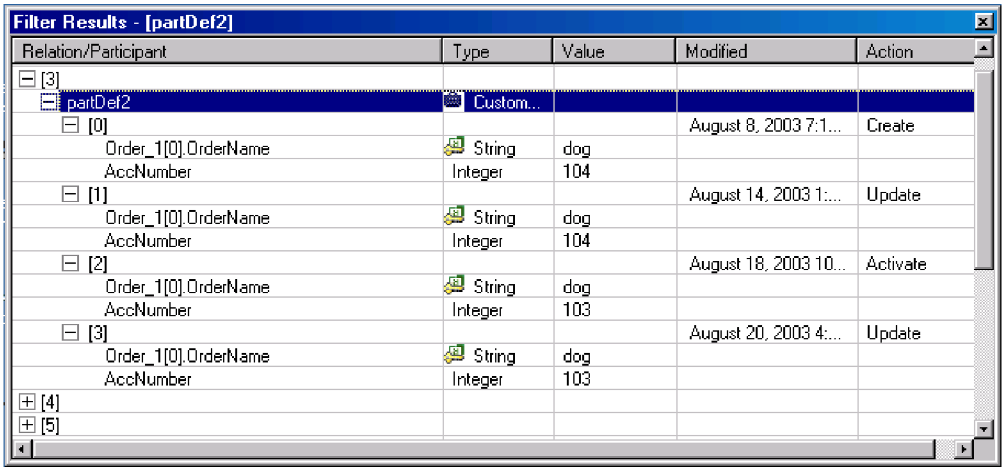

Filter Results - [partDef2]

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

Figure 41. Viewing filtered participant data

Steps for cleaning up participants

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

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

   Relationship Manager displays the “Clean up Participants” dialog box, as shown in Figure 42.

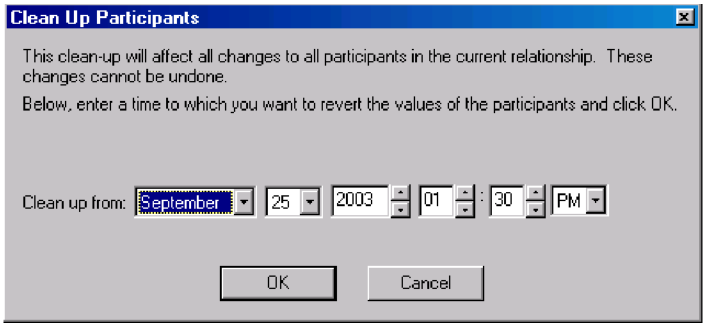

Clean Up Participants

This clean-up will affect all changes to all participants in the current relationship. These changes cannot be undone.

Below, enter a time to which you want to revert the values of the participants and click OK.

Clean up from: September 25 2003 01 : 30 PM

Figure 42. Cleaning up participants

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

   Use the following techniques to enter the date value:

   - Type letters in the text field to cycle through days of the week. For instance, type **S** to cycle through dates that occur on Saturdays and Sundays.

- Click the small up and down arrows to increment or decrement the date by one day.
- Click the large down arrow to display a calendar that you can use to select a date.

4. Click **OK**.

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

Steps for printing relationship data

Relationship Manager allows you to print information about a relationship’s run time data. It creates a tree representation of the run time data, much like the data appears in the tool’s main window. The printing command of Relationship Manager sends the current contents of the relationship tree in the main window to the printer.

Perform the following steps to print relationship run time data:

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

Scheduling jobs

Scheduling jobs allows you to create schedules to manipulate the operational states (start, stop, and pause) of connectors and collaborations. By manipulating component states, you can better manage how InterChange Server processes events. You can distribute the server’s workload over scheduled time periods, thereby reducing traffic and allowing for more efficient resource management. This section covers the following topics:

“Overview of scheduling jobs” on page 99

“Steps for creating schedules” on page 102

“Steps for modifying schedules” on page 103

“Steps for deleting schedules” on page 104

“Steps for displaying schedules” on page 104

“Steps for enabling or disabling schedules” on page 104

Overview of scheduling jobs

Scheduling jobs is done through the Schedule window (see Figure 43). From the Schedule window, you can create, modify, and delete scheduled items. You can see a list of all the schedules that are defined for components, or selectively view schedules based on your requirements. You can also enable or disable all schedules on the server.

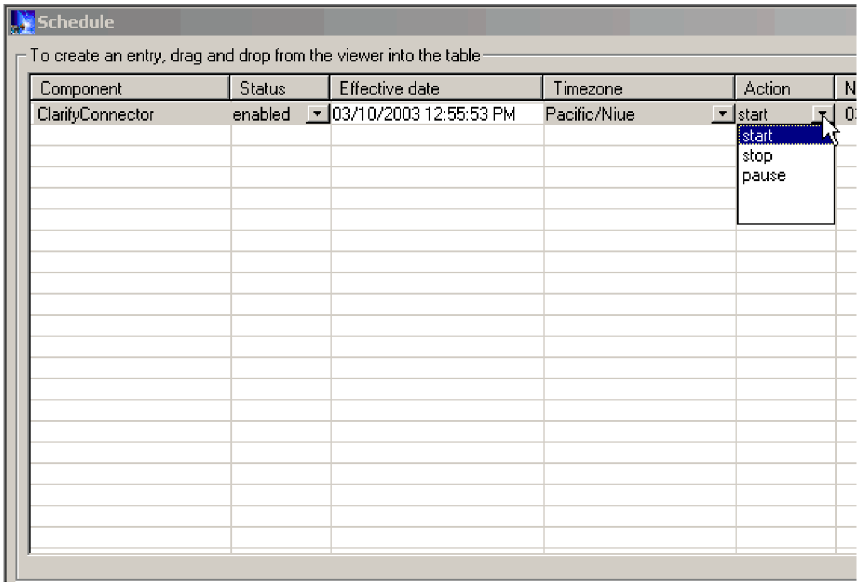
Schedule
To create an entry, drag and drop from the viewer into the table
start
stop
pause

Component	Status	Effective date	Timezone	Action	N
ClarifyConnector	enabled	03/10/2003 12:55:53 PM	Pacific/Niue	start	0:

Figure 43. Schedule window

When you create a schedule for a component, you supply information such as when and how often (recurrence) an action (state change) occurs. By default, no schedules are defined for a component. You can define as many schedules as you want for a component. Once a schedule is set, you can enable or disable its use.

The Schedule window allows you to determine the following items:

Status
Enable turns the schedule on and Disable turns the schedule off. The default status is enabled.

Effective Date
The date and time the schedule is enabled. The default is the current date and time.

Timezone
The time zone where the server is located. The default is Pacific Standard time.

Action
The action the schedule performs. Actions are Start, Pause, and Stop.

Next Occurrence
The next time the scheduled action occurs. If the schedule is non-recurring, the date is the same as the Effective Date. If the schedule is disabled, this field is blank.

Component
The name of the connector or collaboration being scheduled.

Comments
Text field that contains comments you enter about the schedule.

If you choose to make the schedule recurring, you can choose from several options including daily, weekly, or monthly.

Because each schedule consists of one action that occurs at a specified time, to create an interval when the server processes a component, you must define both a schedule to start and end processing. As an example, for a connector, you can create one schedule to start processing events at 1 A.M., and another schedule to pause processing at 3 A.M., daily. Only during that two-hour time period can the connector deliver events to InterChange Server for processing by collaborations that subscribe to that connector.

Overview of scheduling connectors

When you schedule the connector operation, the state you select (start, pause, or stop) determines to what extent work is processed. For example, when you start a connector, it constantly polls an application for new events. When you pause a connector, it stops polling until started again, but is still able to handle service call requests from InterChange Server. A stopped connector is inactive.

By manipulating connector activity with collaboration activity, it is possible to schedule dedicated event processing for an application during a specified time window. To do this, both the collaboration and connector must be running during the same time interval. If the connector was paused, events that were queued can be processed when the connector resumes its activity.

Overview of scheduling collaboration objects

As with connectors, when you schedule the collaboration Object's operation, the state you select (start, pause, or stop) determines to what extent work is processed. To review the collaboration states, see "Viewing collaboration object states" on page 72. For example, when you start a collaboration object, it processes the business objects that it receives from connectors. When you stop a collaboration object, all subsequent events are ignored. So unless you must stop the collaboration object, pause it instead.

Attention: Stopping a collaboration object can cause the connector to delete events as unsubscribed. As a warning, the system produces a message if you select Stop.

When you pause a collaboration object, events remain in the collaboration queue until you restart the collaboration object.

Note: If a scheduled collaboration object is part of a collaboration group, all collaboration objects in that group are scheduled with the same action.

By manipulating collaboration object activity with connector activity, it is possible to schedule dedicated event processing for an application during a specified time window. To do this, both the collaboration object and connector must be running during the same time interval. By assigning different processing windows to collaboration objects that are bound to the same connector, you can distribute the workload, and to some extent, control the amount of traffic a connector must handle. For example, in Figure 44, each collaboration object gets a dedicated time period when the connector is processing only that collaboration object's events.

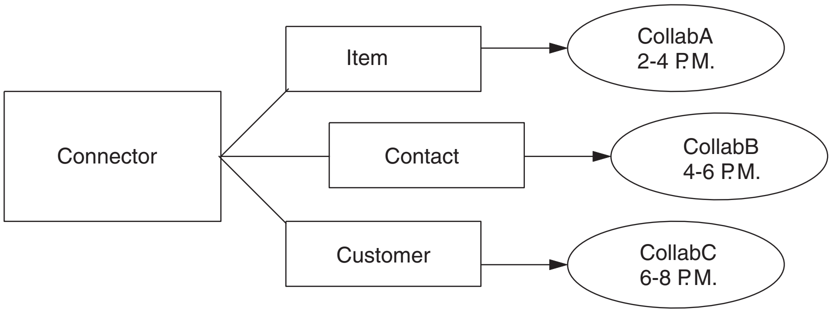

Figure 44. Dedicated processing

Multiple collaboration objects can subscribe to the same business object. In that case, the object is sent to InterChange Server, where it remains until it is picked up by each collaboration object that subscribes to it, when the collaboration object is started after being paused.

Overview of overriding schedules

Using System Monitor, you can override the state of a scheduled component (for example, start a collaboration object that the scheduler stopped a few minutes ago). Or you can set it to a state to one that the scheduler cannot change. For instance, if a collaboration object is scheduled to pause, you can stop it, not allowing the scheduler to pause it (a collaboration object cannot transition from stop to pause). In such a case, the scheduler does not override the manual change, but logs an error instead.

Steps for creating schedules

Perform the following steps to create a schedule for a collaboration or a connector:

1. Open the Schedule window by right-clicking the Schedule folder in System Manager, then selecting ″Edit components' schedule.″ The Schedule window appears, as shown in Figure 43.
2. From System Manager, select the collaboration object or connector to be scheduled and drag it to the Schedule window.

   A new line entry with the name and type of the component is created in the Schedule window (for example, `ClarifyConnector (Connector)`).
3. Enter information about the schedule by clicking the down arrow in each of the schedule cells:
   a. In the Status field, accept **`Enable`** to turn the scheduled item on or select **`Disable`** to turn it off.

      An enabled schedule is effective as soon as you click OK or Apply; a disabled schedule is immediately dormant until enabled. When a schedule is disabled, the Next Occurrence cell is blank to indicate there is no scheduled occurrence for this schedule item.

   b. In the Effective Date field, use the calendar to select the date and time when the scheduled item will occur.

      By default, the current date and time are set. Use the *MM/DD/YYYY hh:mm:ss* format. A 12 or 24 hour clock is used, based on the Time format configured in the Preferences window, which is available from the Edit menu.

   c. In the Timezone field, select the name of the time zone where the scheduled item is being created, if necessary. By default, the time zone for the scheduled item is set to Pacific Standard Time.

For example, the schedule for a connector is created in New York (select **Eastern Standard** time) while InterChange Server is located in Japan. InterChange Server uses this information to determine the local time for the schedule so it can run the job at the appropriate time.

d. In the Action field, select the action to be performed. Actions are **Start**, **Pause**, and **Stop**.

e. Type in any comments you may have in the Comments text cell. A maximum of 255 characters is allowed.

4. If you want this schedule to be ongoing, click the Recurrence check box and enter information about the next occurrence of the action. Click one of the radio buttons to determine a style for inputting the recurrence information and use the down arrow menus to select specific date information:
   - The first radio button, Every, specifies a number and a date element, such as every 2 days or every 3 weeks.
   - The second radio button specifies the date in terms of a monthly event by the day of the week, such as the first Tuesday of every month or the fourth Friday of every month.
   - The third radio button specifies the date as the last day of some number of months, such as the last day of every 3 months.

   If you do not enable the recurrence option, the Next Occurrence field is blank and the schedule expires after it runs. Consistency checks are made to ensure that only one action is scheduled for a particular component on a given date and time. No checks are performed for scheduling conflicts.

   Note: InterChange Server automatically handles changes between standard and daylight savings time for recurring events.
5. Click either of the Show option check boxes to display specific information about schedules. The Show options are:
   - Show Dependencies, which displays schedules for a Collaboration object's bound connectors and collaborations.
   - Show Expired, which displays schedules that have already processed and whose time to run has expired. Only non-recurring schedules expire.
6. Click OK or Apply to create the scheduled item, which is effective once InterChange Server receives the information.

   When InterChange Server and components are geographically distant, there can be a slight delay. If you need to immediate change the state of a component, it is preferable to use System Monitor to start, stop, or pause a component rather than the scheduler.

   Tip: To schedule a time interval when events are processed for a component, you must create a schedule with the Start action and another with the Stop or Pause action. See "Overview of scheduling jobs" on page 99 for information and examples about determining start and end schedules.

Steps for modifying schedules

Perform the following steps to modify an existing schedule for a collaboration or a connector:

1. Right-click the component in System Manager, then select ″Edit Components' schedule.″ The Schedule window appears (see Figure 43).
2. Edit any field in the Schedule list window to change its value.

To edit Recurrence options, click the cursor anywhere on the scheduled item row; the recurrent values for that scheduled item display in the Recurrence pane if they have been assigned.

3. Click OK to save changes and exit, or click Apply to save changes and keep the window open.

Steps for deleting schedules

Perform the following steps to delete an existing schedule for a collaboration object or a connector:

1. Right-click the component from System Manager, then select ″Edit Components’ schedule.″ The Schedule window appears (see Figure 43).
2. Select a scheduled item in the schedule list and click the Delete button (or use the keyboard Delete key) to remove the schedule.
3. Click OK to save changes and exit, or click Apply to save changes and keep the window open.

Steps for displaying schedules

Perform the following steps to display a schedule or a group of schedules:

1. Select and open an object for displaying schedules:
   - Collaboration or connector icon from System Manager. The Schedule window displays all schedules defined for that object. If Show Dependencies is active, all schedules for components connected to the object are displayed.
   - Collaboration or connector folder from System Manager. The Schedule window displays schedules for all objects in that folder.
   - InterChange Server from System Manager. The Schedule window displays all schedules defined for that server, with times shown for the server’s time zone.
   - If you choose the Schedule option from the main window, all schedules in the system are displayed.
2. Click any of the column headings to sort schedules by that column.

Steps for enabling or disabling schedules

Perform the following steps to selectively disable or enable schedules:

1. Select an object for displaying schedules.

   See “Steps for displaying schedules” on page 104.
2. Enable or disable the schedule:
   - To enable or disable all schedules, click either the Enable All or Disable All radio button.
   - To enable or disable a single schedule, click the down arrow in the Status column and choose the Enable or Disable option.
3. Click Apply to complete this task.
4. Click OK to exit.

Backing up system components

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

Note: Backed up repositories from previous versions should not be restored in the current version, since there is no guarantee of CWBF format enforcement in Bi-Directional (BiDi) enabled releases. This could be due to several reasons not related to the new release and the subsequent BiDi support, for example, non-BiDi enabled adapters or the custom codes associated with system components not in CWBF. For more information about BiDi, refer to the *Technical Introduction to IBM WebSphere InterChange Server* and the *Business Object Development Guide*.

This section covers the following topics:

"Backup schedule planning" on page 105

"Component backups" on page 106

Backup schedule planning

It is important for you to plan and carry out procedures for regularly scheduled backups of the IBM WebSphere InterChange Server system. The more frequently you perform backups, the less data you need to recover in the event of data loss.

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

- Static data rarely changes and should be backed up only when changed. For example, static configuration data stored in the IBM WebSphere InterChange Server repository needs backing up only when it has changed. Static data should also be backed up before any planned reinstallations or upgrades to the system. Following is a partial list of static data in the IBM WebSphere InterChange Server system:
  - IBM WebSphere InterChange Server repository (except for relationship tables)
  - Custom collaborations components, such as Java class files (`.class`), and message files (`.msg`)
  - Custom connectors
  - Map components, including: map design files and Java class files (`.class`).
- Dynamic data constantly changes and should be backed up on a regular basis. For example, the relationship tables maintain the instance data for relationship definitions. Since relationship instance data is maintained continuously, regularly back up this data as well as application data.

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

  Following is a partial list of dynamic data in the IBM WebSphere InterChange Server system:
  - IBM WebSphere InterChange Server cross-reference database
  - IBM WebSphere InterChange Server relationship tables
  - IBM WebSphere InterChange Server WIP (event management) and transaction tables
  - WebSphere MQ queue data

- IBM WebSphere InterChange Server connector archive tables (this is part of the application backup; all events since the last backup should be archived)
- Log files (as desired for historical information)

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

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

Component backups

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

"Relationship table backups" on page 106

"Repository backups" on page 107

"System installation file backups" on page 107

"Collaboration class file backups" on page 107

"Archive table backups" on page 107

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

Relationship table backups

Relationship tables are backed up using the standard backup utility for the database where these tables reside. Schedule this backup to coincide with the corresponding application backups. If you back up applications at different times, back up the relationship tables each time you back up an application. There are often static relationship tables within the relationship database. Although this data is static, it is recommended that you back up *all* relationship tables together. Make sure the IBM WebSphere InterChange Server system is in a quiescent state when backing up the relationship tables. For more information on bringing the system to a quiescent state, see "Steps for shutting down InterChange Server" on page 50.

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

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

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

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

Repository backups

Repository tables are backed up using the `repos_copy` command. For more information on this command, see "Using repos_copy" on page 108. Back up the repository whenever it is modified and before and after performing a reinstallation or an IBM WebSphere InterChange Server software upgrade. The IBM WebSphere InterChange Server system does not need to be in a quiescent state when backing up the repository.

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

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

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

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

System installation file backups

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

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

Collaboration class file backups

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

Archive table backups

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

Using repos_copy

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

`Repos_copy` is also used to migrate components from earlier versions to the current release. If you are working with old-format components, first migrate your components. `Repos_copy` does not support "-ar, -arp, -vr, -vp -xCompilePackage", and has limitations for "-o", and all the -xCompile options when working with old-format components.

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

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

Note: `Repos_copy` is not BiDi enabled. If you attempt to use `repos_copy` to restore a repository that includes BiDi data that is not in CWBF format, the data will be in an inconsistent state. For more information about BiDi, refer to the *Technical Introduction to IBM WebSphere InterChange Server* and the *Business Object Development Guide*.

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

This chapter has the following sections:

- "Repos_copy syntax" on page 108
- "Repos_copy usage scenarios" on page 114
- "Locale for repos_copy files" on page 119

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

Repos_copy syntax

Table 20 on page 109 describes the options of repos_copy and their arguments, and shows the correct case usage for the options and the lack of spacing between the option and its argument. The syntax shows that the options between curly braces ({}) represent a set of options that are required. If you do not specify the `-u`, `-p`, `-i`, `-o`, or `-s` options at the command line, then `repos_copy` prompts you for them. If you do not specify them when prompted, repos_copy does not execute. Options enclosed in brackets ([]) are optional.

Note: Some new arguments have been added in release 4.2, and some arguments from the previous release have been removed. For a list of these arguments, see “New arguments in release 4.2” on page 113 and “Arguments removed in release 4.2” on page 113.

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

Table 20. Repos_copy command options

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

Table 20. Repos_copy command options (continued)

Option	**Description**
-do*EntityType:Entity[+EntityType:Entity2]*	Specifies the entities to be deleted from the repository. See Table 21 on page 113 for the list of entity types and keywords. If the object has no referents—other components that depend on it—then the deletion takes place. If the object has referents, then the deletion fails and a message is displayed. The behavior is the same in both design mode and production mode. For more information about starting the server in design mode or production mode, see the *System Implementation Guide*.
-e*EntityType:Entity1[+EntityType:Entity2...]*	Exports one or more referenced first-class entities. A first-class entity is a business object, collaboration object, collaboration template, connector, database connection pool, map, or relationship. You identify the entity to load or unload by specifying one of the keywords in Table 21 on page 113. Follow the *EntityType* keyword with a colon (:) and the name of the entity. Use the "+" to specify more than one entity. When combined with the -o option, the -e option unloads the data to an output file.
-f*entityFile*	This option is similar to the -e option except that the names of the entities to be imported are stored in a file. The file should contain references to the entities, with the following conditions: • The entity names must follow after the proper entity type keyword. The entity types and their keywords are listed in Table 21 on page 113. • A colon must separate the entity type from the entity name. • There must be a new line separating each entity reference. When combined with the -o option, this option exports the components to a package.
-i*filename*	Deploys the specified package file to the repository. If you omit the input file name value, the command interactively prompts you to enter the name of the input file. The file can be either a .jar file containing objects in XML format, or a file in text format from a release prior to 4.2.0. The .jar files created by repos_copy or System Manager have a particular structure which must be maintained for any subsequent imports of such a file to be successful. You should not, therefore, ever modify an input file manually.

Table 20. Repos_copy command options (continued)

Option	Description
-k	Overrides the default behavior of repos_copy when it finds a Mercator map in the package file it is loading. By default, repos_copy exits if it encounters a Mercator map. If you use the -k option, repos_copy skips over any Mercator maps in the package file and proceeds with the deployment process.
-mode	Returns the mode of the server. For more information about InterChange Server modes, see the *Implementation Guide for WebSphere InterChange Server*.
-nc*encoding*	Specifies the character encoding when importing a text-based repository file fromreleases prior to 4.2.0. To ensure successful data, you must use encoding that is inconsistent with text-based repositories. Failure to do this may result in corrupt data. For a list of valid character encodings, see the Java documentation about the String class.
-o*outfilename*	Exports the components in the repository to the specified package file. You must specify the name of the package file. If the file already exists then repos_copy prompts you to overwrite it or not. The output file is in .jar format, and contains the component definitions in XML format, as well as .java source files for components that have them. This option cannot be combined with the -i or -d options, nor can it export components in text format as it did in previous releases. Repos_copy does not append the .jar extension, so you must specify it when specifying the name of the output file.
-p*password*	Specifies the password for the user name supplied with the -u option. The password case-sensitive. If you do not specify this option then repos_copy prompts you for the password.
-r*	This option is similar to the -r option; it allows you to import relationship definitions and not create the run time schemas for any of them.
-r*relationshipName1[:relationshipName2]*	Loads the named relationship definition(s) into the repository without creating its run time schema.
-s*serverName*	Specifies the name of the InterChange Server instance with which repos_copy should interface. The name is case-sensitive. If the server name is not specified, the tool prompts for a server name.
-summary	This option prints a list of components in the server repository (they are identified as “artifacts” rather than as components in the output). The output is in XML format. You can combine this option with the -o option to print the output to a file rather than the console.

Table 20. Repos_copy command options (continued)

Option	Description
`-u`*`username`*	Specifies the user name to log in to InterChange Server. If no user name is specified, repos_copy prompts for a user name.
`-v`	Prints the version number of the program that the repos_copy utility executes.
`-vp`	This option validates a package file. The server validates packages against the repository and makes sure that the dependencies among the components in the package are resolved. If the validation is not successful, repos_copy prints a list of the missing dependencies. This option does not make any changes to the repository; it just validates the package file. When using the `-vp` option you must also use the `-i` option to specify the package file to be validated.
`-vr`	This option validates the repository. The output message indicates whether the validation is successful or not. If the validation is not successful, repos_copy prints a list of the missing dependencies.
`-wi`	When this option is specified, repos_copy does not display any warnings that occur during the compilation of collaboration templates or maps. Only errors that occur during compilation are displayed. This allows the user to ignore warnings about deprecated methods, for instance.
`-xCompileAll`	Compiles all collaboration templates and maps in the repository. Valid only for collaboration templates and maps created using release 4.2 or later.
`-xCompileAllCollabs`	Compiles all collaboration templates in the repository. Valid only for templates created using release 4.2 or later.
`-xCompileAllMaps`	Compiles all maps in the repository. Valid only for maps created using release 4.2 or later.
`-xCompileCollab:`*`collabTemplateName[+collabTemplateName]`*	Compiles the specified collaboration templates in the repository. Valid only for templates created using release 4.2 or later.
`-xCompileMap:`*`nativeMapName[+nativeMapName]`*	Compiles the specified maps in the repository. Valid only for maps created using release 4.2 or later.
`-xCompilePackage`	This option automatically compiles the package being deployed to the server. Since the production-mode server automatically compiles all packages, this option applies only to design-mode servers. For a full description of InterChange Server modes, see the System Implementation Guide. **Note:** This option works only if you are deploying components from release 4.2. If the components are from a prior release, this option will be ignored.

Table 20. Repos_copy command options (continued)

Option	Description
`-xdi`	For a full description of InterChange Server modes, see the System Implementation Guide.
`-xdn`	For a full description of InterChange Server modes, see the System Implementation Guide.
`-xmsp`	This option imports and exports membership and security information, allowing you to upgrade without having to recreate Roles and Security Policy. For a full description of InterChange Server modes, see the System Implementation Guide.

Table 21. Keywords for different entity types

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

New arguments in release 4.2

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

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

Arguments removed in release 4.2

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

Table 22. Arguments removed from repos copy

Arguments removed	Reason for removal

`[-xCompileUpdated]`
`[-xCompileUpdatedCollabs]`
`[-xCompileUpdatedMaps]`
	All compile update options have been removed, because the server does not support maps or templates prior to release 4.x.

`[-xUncompress]`	In the new package format, all definitions are stored in the Java Archive (JAR) format, instead of with a proprietary compression algorithm. The `-xUncompress` argument is therefore no longer necessary.

Table 22. Arguments removed from repos copy (continued)

Arguments removed	Reason for removal
[-eProject][-w]	Project type is no longer supported for the -e option. Projects are now maintained in the local file by System Manager, instead of being maintained by the server.

Repos_copy usage scenarios

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

- "Example of printing the repos_copy command"
- "Example of validating a package"
- "Example of validating a package"
- "Example of deploying a package to the repository" on page 115
- "Example of validating the repository" on page 116
- "Example of deleting components from the repository" on page 117
- "Example of exporting components to a package" on page 118
- "Example of printing a list of components in the repository" on page 119

Example of printing the repos_copy command

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

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

Example of validating a package

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

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

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

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

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

Example of deploying a package to the repository

The -i option allows you to deploy a package of components to the repository. If you do not specify the name of the package file then you are prompted to enter it.

The following example shows a a file named WebSphereICS420DEVServer.jar being deployed to a repository:

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

Example of duplicate components: Commonly there will be components with the same name in the package file as there are in the repository. In this case you must decide whether or not you want to replace the components in the repository with those in the package file. The -ai option specifies that duplicate components should not be loaded into the repository:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-iCustomer.jar -ai
```

If you want to replace all the duplicate components in the repository, use the -ar option as in the following example:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-iCustomerSyncInterface.jar -ar
```

You can use the -arp option to interactively replace duplicate components in the repository. This lets you decide for each individual duplicate component whether it should be replaced or not.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-iCustomerSyncInterface.jar -arp
```

Note: The -ar and -arp options only work with release 4.2.0 or later.

Example of compiling and creating schemas: For maps and collaboration to execute at run time, the maps and collaboration templates defined in the repository must be compiled. For relationships to function properly at run time, their schemas must be created.

When you deploy components to a server running in production mode, all templates are automatically compiled and all relationship schemas are created. For the deployment to succeed, then, the code of the map and collaboration templates must be valid and InterChange Server must be able to communicate with the databases specified in the settings of the relationship definitions.

When you deploy components to a server running in design mode, the templates are not automatically compiled; relationship schemas are automatically created. There are options you can use to compile the templates, however, and there are options to not create relationship schemas.

The following example uses the -xCompilePackage option and does not use any form of the -r option. The result is that when the package specified by the -i option is deployed, the maps and collaboration templates are compiled and schemas are created for the relationships:

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

You may not want relationship schemas created when you do a deployment. For instance, if you are deploying a package from one environment to another and did not change the properties of the relationships to use the database resources in the new environment then you will not want the schemas created until after you have changed the relevant properties. The following example uses the -r* option to not create schemas for all of the relationships in the package being deployed:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-iWebSphereICS420DEVServer.jar -xCompilePackage -r*
```

Note: You can use the -r option without the asterisk to specify the names of individual relationships whose schemas should not be created. For instance, -rCustomer:Order would not create schemas for the Customer and Order relationships, but would still create schemas for any other relationships in the package being deployed.

Important: Although there are options to compile maps and collaboration templates after deployment, there is no way to either through repos_copy or System Manager to create the schema for a relationship other than during deployment. So, if you chose not to create the schema for a relationship during deployment because you needed to change the database settings, then you need to re-deploy the relationship afterwards and allow repos_copy to create the schema for the relationship.

Example of validating the repository

The repository must be in a valid state for a server instance to start in production mode. The reason for this is that ultimately the repository must be valid for the server to process flows successfully. Use the -vr option to validate a server repository, as in the example below:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull -vr
```

If the server is valid then repos_copy writes the following output to the console:

```
Validation Succeeded.All Dependencies Resolved.
```

If the repository is not valid then repos_copy prints a list of the dependencies that must be resolved.

Example of compiling components in the repository

If you deployed maps or collaboration templates to the repository and did not compile them during deployment, you can use repos_copy to compile them afterwards. This can be useful in situations where there are many components to deploy because deployment can take a long time and compiling can make the operation take even longer. Waiting until after the deployment has succeeded to do the compilation task can reduce the risk of spending an even greater amount of time migrating the environment if an error occurs.

The following example shows the use of the -xCompileAll option to compile all maps and collaboration templates in the the repository:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-xCompileAll
```

There are options to compile all of either type of component as well. Use -xCompileAllCollabs to compile all the collaboration templates, and -xCompileAllMaps to compile all the maps. The example below shows the use of -xCompileAllMaps:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-xCompileAllMaps
```

Just as you can compile all of one type of component, you can also compile an individual component. Use the -xCompileCollab or -xCompileMap option followed by a colon and the name of the collaboration template or map to compile a single component. The example below would compile a collaboration template named CustomerSync:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-xCompileCollab:CustomerSync
```

Note: The compilation options only work with WebSphere InterChange Server version 4.2.0 and later.

Example of deleting components from the repository

There are several options provided by repos_copy for deleting components in the repository. You can delete the entire repository, individual components, and individual components as well as any components that reference them.

Note: Components must be inactive for you to delete them. If you delete a single component then you must deactivate it first or the delete operation will fail. If you want to delete a component and all the components that reference it, you must deactivate not only the single component, but all those that reference it as well. You can delete the entire repository while the components are in an active state. Use System Monitor or web-based System Monitor to manage the states of components. System Monitor and web-based System Monitor are described in the *System Administration Guide*.

Example of deleting the entire repository: Use the -d option to delete all of the components in the repository. The following example shows the syntax:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin
-pnull -d
```

Repos_copy presents a prompt asking if you want to delete the entire repository or not.

Example of deleting components without referents: If a component does not have any referents—other components that reference it and require it to exist in order to perform their function in the system—then you can delete the individual component.

Use the -do option followed by the entity type, a colon, and the name of the component. The entity types are listed in Table 21 on page 113. The following example deletes the relationship named **Customer**:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin
-pnull -doRelationship:Customer
```

Example of deleting components with referents: If a component does have referents—other components that reference it and require it to exist in order to perform their function in the system—then you can only delete the component if the server is running in design-mode, and by using certain options.

Forcing a delete in spite of references: If a component has referents, repos_copy will not let you delete it with the -do option. You must use the -dfo option to force deletion of a component with referents. Forcing deletion of a component that has referents will leave the repository in an inconsistent state, and a server running in production mode does not permit that, so this option only works with a design-mode server. The following example shows the use of the -dfo option to delete the Order business object in spite of the fact that other components in the system (such as maps and relationships) have references to it:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-dfoBusObj:Order
```

Deleting the referents as well: Another way you can delete a component that has referents is to use the -deep option to delete the referents as well. This deletes the component and all of the components that have references to it. The following example shows the use of the -deep option when using the -do option to delete the Customer business object:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-doBusObj:Customer -deep
```

This option, unlike the -dfo option, is supported with servers running in production mode because the deletion of the referents along with the component guarantees that the repository remains valid. Keep in mind, however, that it can result in many components being deleted; you should be aware of the implications of this action prior to taking it.

Example of exporting components to a package

The -o option allows you to export components from the repository to a package. You must specify the name of the package file. When the -o option is used alone the entire repository is exported to a file, as in the following example:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-oWebSphereICS420DEVServer.jar
```

You can specify individual components to be exported by using the -e option. You must use the -e option with the appropriate EntityType keyword listed in Table 21 on page 113, and must follow the keyword with the name of the component. You can specify multiple components by concatenating them with the plus (+) sign. In the following example, the Customer business object and CustomerSync collaboration template are exported to a package named CustomerSyncInterface.jar.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-eBusObj:Customer+CollabTemplate:CustomerSync -oCustomerSyncInterface.jar
```

You can use the -deep option to export the dependencies of a component as well. In the previous example, the Customer business object was exported, but none of its child business objects were. The following example uses the -deep option to export the CustomerSync_ClarifyToSAP collaboration object and all of its dependencies.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-eCollaboration:CustomerSync_ClarifyToSAP -oCustomerSyncInterface.jar -deep
```

If you want to export specific components, but do not want to have to enter the entity type keyword and component names, you can store them in a text file and use the -f option. This is very convenient when you want to frequently export the same components. The following example uses the -f option to load the components listed in a text file named Components.txt :

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull
-fComponents.txt -oCustomerSyncInterface.jar -deep
```

The contents of the file Components.txt are shown below; a paragraph return follows each entity type keyword and name combination:

```
BusObj:Customer
Relationship:Customer
CollabTemplate:CustomerSync
```

Note: Repos_copy and System Manager are unfortunately inconsistent with respect to what they identify as “dependencies”. If you attempt to delete a component using repos_copy but there are components that depend upon it then repos_copy lists those referring components as dependencies. However, if you right-click the component in System Manager and select **Show Dependencies** from the context menu the tool lists the components that the selected component depends on.

Example of printing a list of components in the repository

You can use the -summary argument when executing repos_copy to print a list of the components in the repository. The output is presented in XML format. Although it is not particularly useful to view at the command line, you can combine the -summary argument with the the -o argument to redirect the output to a file and then view the file in a browser or XML editor. The command usage in this case would be the following:

```
C:\>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull -summary -oRepository.xml
```

Locale for repos_copy files

The repos_copy utility reads metadata from the repository and writes the data out to files in Unicode (UTF-8 format). It also reads such files and loads them into the repository in Unicode (UTF-8 or UCS-2, as the underlying repository database dictates).

Repos_copy files created with IBM WebSphere InterChange Server version levels earlier than 4.1.1 can be loaded ino the repository correctly only if the dates and times for the component schedules are in full US format. (This is usually not an issue. Repos_copy saves all schedule dates in full US format only. The incompatibility could typically arise if the repos_copy files have been manually edited.)

Administering end-to-end privacy

The security of messages and business data on a system is critical, from the moment they leave a source adapter, through their journey into the InterChange Server, right up until they reach a destination adapter. Critical to any secure system is end point verification. The IBM WebSphere InterChange Server provides security at each end point of the information flow, ensuring that your information is secure from end-to-end.

Most business communication in InterChange Server is transported over asynchronous systems such as JMS and MQ Series, causing messages to be stored on disk at the queue manager while they wait for processing. End-to-end privacy ensures that these messages are secured at this level.

In order to use end-to-end privacy to protect your messages, you must activate it in the appropriate configuration file. End-to-end privacy can be turned on or off for each individual adapter.

It is important to note that the configuration of end-to-end privacy using System Manager will only effect the messages from InterChange Server tothe adapter, while the configuration of the adapter using Connection Configurator will only effect the message from the adapter to InterChange Server.

Note: For in-depth information on end-to-end privacy concepts and functionality, refer to the *Technical Introduction to IBM WebSphere InterChange Server*.

This section covers the following topics:

"Steps for activating end-to-end privacy using System Manager" on page 120

"Steps for activating end-to-end privacy using Connection Configurator" on page 121

"Steps for changing the privacy configuration using System Manager" on page 121

"Steps for changing the privacy configuration using Connection Configurator" on page 121

"Administering keys and keystores" on page 121

Steps for activating end-to-end privacy using System Manager

Perform the following steps to activate end-to-end privacy using System Manager:

1. On the Privacy tab, enter the path to the keystore. For additional information on keystores, see "Administering keys and keystores" on page 121.
2. Enter the password for the keystore.
3. To import a specific privacy setting, select the Import Privacy Setting button and select one of the available connectors. This loads the privacy configuration specified by the specific connector's configuration file.
4. To set a general privacy setting, select a message type from the drop down list. Available choices are:
   - All
   - Admin
   - BO
5. Select a security level from the drop down list. Available choices are:
   - None
   - Privacy
   - Integrity
   - Integrity plus Privacy
6. Select a destination for the messages, for example, System Test Connector or Destination Connector. This is an optional distinction used only when end-to-end privacy has been set for the specific connector.
7. To set a privacy setting for an individual business object, enter the name of the business object or select a business object from the available list.
8. Select a security level from the drop down list. Available choices are:
   - None

- Privacy
- Integrity
- Integrity plus Privacy

9. Select a destination for the messages, for example, System Test Connector or Destination Connector. This is an optional distinction used only when end-to-end privacy has been set for the specific connector.

Steps for activating end-to-end privacy using Connection Configurator

Perform the following steps to activate end-to-end privacy:

1. On the Connector Configurator tab, select the Support tab.
2. From the listing, select the drop down list under the Privacy heading to assign the appropriate privacy level for each individual business object.Available choices are:
   - None
   - Privacy
   - Integrity
   - Integrity plus Privacy
3. Save your connection configuration to activate the privacy settings.

Steps for changing the privacy configuration using System Manager

Perform the following steps to change the end-to-end privacy parameter using System Manager:

1. On the Privacy tab, update the following information:
   - Keystore path
   - Keystore password
   - General privacy settings
   - Individual business object privacy settings

Steps for changing the privacy configuration using Connection Configurator

Perform the following steps to change the end-to-end privacy parameter using Connection Configurator:

1. On the Security tab, update the privacy setting for any individual business object.
2. Save your connector configuration.

Administering keys and keystores

A keystore is a password protected file used to securely store the public and private keys used for privacy verification. A keystore is present for the server, as well as for each individual adapter. The IBM WebSphere InterChange Server contains an InterChange Server private and public keystore, as well as the public keys of each adapter. Each individual adapter keystore contains the adapter private and public keystore, as well as the public key of the InterChange Server.

The full path to the keystore and the applicable password, which is encrypted during startup, is contained in the configuration file. The password for the private key should be identical to the password for the keystore.

Steps for creating keys using the connector configurator

Perform the following steps to create the keys and keystores using the graphical unterface provided with the connector configurator:

1. On the Connector Configurator screen, select the Security tab. From this tab, you can turn on end-to-end privacy, set your privacy levels, maintain keys and set adapter access control.
2. Select the Generate Keys button. The Generate Keys screen displays.
3. Enter the following information to complete the key building process:
   - Certificate association
   - Generation algorithm
   - Output keystore
   - Keystore password
   - Private key password
   - Any additional key options
4. Enter one of the following pieces of information. Only one is required to build the key:
   - Common name for the key
   - Organization unit
   - Organization name
   - Locality name
   - State name
   - Country name
5. Select the OK button to save the key information.

Steps for creating keys using the keytool

Perform the following steps to create the keys and keystores using the keytool:

1. Open the keytool found in the JDK_HOME/bin directory.
2. Create a public and private key entry for the server by entering the following command line, where `name` equals the keystore name, `password` equals the keystore password and `IC.keystore` equals the keystore file name:
   ```
   keytool -genkey -alias name -keyalg RSA -keypass password -storepass
   password -keystore IC.keystore
   ```
3. Export the public key of the server to a file by entering the following command line:
   ```
   keytool -export -alias name -storepass password -file IC.cer -keystore
   IC.keystore
   ```

   Note: The adapter agent will import the IC.cer file when it imports the public key of the server into it's keystore.
4. Create a public and private key entry for the adapter agent by entering the following command line, where `connectorname` equals the keystore name, `password` equals the keystore password and `Adapter.keystore` equals the keystore file name:
   ```
   keytool -genkey -alias connectorname -keyalg RSA -keypass password
   -storepass password -keystore Adapter.keystore
   ```
5. Export the public key of the adapter to a file by entering the following command line:
   ```
   keytool -export -alias connectorname -storepass password -file
   Adapter.cer -keystore Adapter.keystore
   ```

6. Turn on privacy settings by importing the key for the adapter agent into the server keystore by entering the following command line:
```
keytool -import -v -trustcacerts - alias connectorname -storepass
password -file Adapter.cer -keystore IC.keystore
```
7. Import the server's public key into the adapter agent's keystore by entering the following command line:
```
keytool -import -v -trustcacerts - alias connectorname.queue.manager
-storepass password -file IC.cer -keystore Adapter.keystore
```

Steps for exporting the adapter public key

Perform the following steps to export the adapter public key:

1. On the Connector Configurator screen, select the Security tab.
2. Select the Export Adapter Public Key button.
3. On the Export Adapter Public Key screen, enter the following information:
   - Output certificate
   - Input keystore
   - Keystore password
   - Certificate association
   - Any additional key options
4. Select the Ok button to export the adapter key.

Steps for importing the server public key

Perform the following steps to import the server public key:

1. On the Connector Configurator screen, select the Security tab.
2. Select the Import Server Public Key button.
3. On the Import Server Public Key screen, enter the following information:
   - Output keystore
   - Input certificate
   - Keystore password
   - Private key password
   - Certificate association
   - Any additional key options
4. Select the Ok button to import the server key.

Administering role-based access control (RBAC)

One of the key features to the IBM WebSphere InterChange Server is the ability to authorize permissions for users accessing the system using roles, known as Role-based access control (RBAC). Roles can easily be defined by the Administrator and assigned to a group of users, restricting access to key components only to verified users. Roles can be assigned along functional associations and greatly reduce the administrative burden. Assigning a role to a user or users allows them to access only the components of the system included in the role definition.

Use of RBAC functionality ensures that only an Administrator, or users with permission to administer roles, would be allowed to create users and assign roles. If RBAC is not active on the server, any user can create users and roles with no verification.

Note: When you activate RBAC in InterChange Server, the RBAC run time status displays on the System Manager screen.

For information on configuring Role-based access control, see "Steps for configuring RBAC security" on page 39. For in-depth information on Role-based access contol concepts and functionality, refer to the *Technical Introduction to IBM WebSphere InterChange Server*.

Note: The Failed Events Manager uses RBAC functionality to establish roles which administer access control to failed events information. For more information on the Failed Events Manager, refer to the *Problem Determination Guide*.

This section covers the following topics:

"Steps for setting up RBAC"

"Steps for deactivating RBAC" on page 125

"Administering roles" on page 125

"Administering users" on page 126

"Administering user and role assignments" on page 127

"Administering security policy permissions" on page 127

"Administering membership and security policy information" on page 128

"Administering the RBAC password" on page 129

"Security Administration" on page 129

Steps for setting up RBAC

Before setting up RBAC, at least one user must be assigned the role of Administrator. If no user is assigned an Administrator role, the server will always re-boot with RBAC disabled. Perform the following steps to set up role-based access control:

1. On the Security-RBAC tab, select the check box for Enable RBAC.
2. Select the user registry to which to apply role-based access controls, that is, Repository or LDAP.

   Note: If you select the LDAP user registry, you must ensure that the server privacy keystore is set up in order to assure correct functioning.
3. In the Server Start User Name field, enter the user name to start the server.
4. In the Server Start Password field, enter the password associated with the username.
5. If you selected Repository, enter details in the following fields:
   - Host name
   - Database
   - Port Number
   - User Name
   - Password
   - Max Connections, which is the maximum number of connections that the user can open

- Max connect retries, which is the maximum number of times you can attempt to start a connection
- Connect retry interval, which is the amount of time between connection retries

6. If you selected LDAP, enter details in the following fields:
   - LDAP Url, which is the url of the LDAP installation
   - Username, which is the user account and is not case-sensitive
   - Password, which is the password for the user account
   - Userbase DN, which is the base distinguished name and acts as the root of all searches and updates
   - Username attribute, which the attribute in the schema that InterChange Server uses as a username
   - Search criteria, which is the search criteria to use when retrieving LDAP users and is optional
   - Max search returns, which is the maximum number of entries returned from a search
   - SSL, which when set to `True` secures the connection using SSL protocol
7. To turn on Audit settings, select the check box for Enable Audit and enter details in the following fields:
   - Audit log directory, which is the path of the audit log file
   - Audit log frequency, for example, Daily, Weekly or Monthly
   - Audit file size, which is the maximum size for the audit file in MB

Steps for deactivating RBAC

Perform the following steps to deactivate RBAC:

1. On the Security-RBAC tab, select the check box for Enable RBAC. Disabling RBAC functionality causes all the fields in the display to become grayed.

Administering roles

Role-based access control (RBAC) supports multiple users and enhanced security features based on roles. A role is a collection of users who share common functionality. Assigning functions into roles allows the administrator to work more effectively by reducing the burden on the administrator during the assignment of permissions.

If a role is no longer necessary for the functioning of the server, you may choose to delete that role from the listing. Once a role is deleted, all role references are removed from the applicable users.

Note: The Failed Events Manager also uses RBAC functionality to establish roles which administer access control to failed events information. For more information on the Failed Events Manager, refer to the *Problem Determination Guide*.

Steps for creating roles

Perform the following steps to create a role:

1. On the Context Menu, select New Role. This displays the Role Name dialog box.
2. Enter the role name. Once you name a role, it cannot be renamed.
3. Enter a role description, if necessary. Role description is an optional field.

Steps for deleting roles

Perform the following steps to delete a role:

Note: The role `administrator`is the default and cannot be deleted. It is case-sensitive.

1. On the Context Menu, select Delete Role.
2. Select the role name. Once you delete a role, it cannot be restored.

Administering users

On the User and Roles Management screen, roles are listed downward in a tree directory display. You can assign a user to any number of roles. Users assigned to a role are listed in the tree directory beneath the role to which they are assigned, making for a quick and easy scan of permissions and responsibilities.

Additionally, you can import or export user information for use with the RBAC functionality.

Steps for adding users

Perform the following steps to add users to RBAC:

1. On the Context Menu, select New User. This displays the New User dialog box.
2. In the Username field, enter the name of the user.
3. In the Password field, enter the password for the user.

Steps for deleting users

Perform the following steps to delete users from RBAC:

Note: `Guest` is the only default user and cannot be deleted.

1. On the Context Menu, select Delete User.
2. Select the user name. This removes the user from all pre-assigned roles.

Steps for Importing users and passwords

Perform the following steps to import users and passwords into RBAC:

Note: When DATABASE is the user registry, support is available for importing users. However, this function is not supported for the LDAP user registry. It is recommended that you create a central user registry database or central LDAP registry, enabling multiple InterChange Server machines to use this central repository as opposed to transfering the user registry across various InterChange Server Machines.

1. On the Context Menu, select Import >> User Registry. This displays the Import dialog box, where you specify the path for the binary file. This path should be valid on the server machine which is running the InterChange Server.
2. Select the file to import.

Steps for exporting users and passwords

Perform the following steps to export users and passwords into RBAC:

Note: When DATABASE is the user registry, support is available for exporting users. However, this function is not supported for the LDAP user registry. It is recommended that you create a central user registry database or central LDAP registry, enabling multiple InterChange Server machines to use this central repository as opposed to transfering the user registry across various InterChange Server Machines.

1. On the Context Menu, select Export >> User Registry. This displays the Export dialog box, where you can specify the file path.
2. Select the destination for the file to export. This path should be valid on the server machine which is running the InterChange Server.

Administering user and role assignments

Assigning roles to the available users greatly reduces the burden upon the administrator to assign individual permissions to vital functionality. Users can be assigned to numerous roles, all regulated by the user's login ID. Users assigned to a role are listed in the tree directory beneath the role to which they are assigned. Perform the following steps to assign roles to users:

Steps for assigning roles to users

Perform the following steps to assign roles to users:

1. On the Context Menu, select the user to which you want to assign roles.
2. Select Add Role. This displays the Add Role dialog box, which lists all available roles.
3. Select single or multiple roles to assign to the user. This lists the assigned users under the roles display.

Steps for removing users from roles

Perform the following steps to remove users from the roles listing:

1. On the Context Menu, select the user you want to remove from the role permissions.
2. Select Remove Role. This removes the user from the role listing and removes all role permissions from the user profile.

Administering security policy permissions

As an administrator, you can assign permissions to default roles within RBAC. These security policies are listed in a tree directory, along with the operations that each role is allowed to access.

Table 23lists the operations that can be secured in a server.

Table 23. Secured Server Operations

Secureable component	Access-controlled operations
Server	1. Start 2. Shut Down 3. Security/Administering users/Roles 4. Monitoring 5. View Failed Events 6. Deploy 7. Export 8. Delete 9. Compile 10. Export config files 11. Deploy config files
Collaboration Templates	1. Compile

Table 23. Secured Server Operations (continued)

Secureable component	**Access-controlled operations**
Collaboration Objects	1. Start 2. Stop 3. Pause 4. Shutdown 5. Execute (AccessFramework call) 6. Resolve transactioanl status 7. Submit Failed events 8. Delete Failed events 9. Cancel LLBP flow
Connectors	1. Start 2. Stop 3. Pause 4. ShutDown Agent 5. Submit Failed Events 6. Delete Failed Events
Business Objects	
Maps	1. Compile 2. Start 3. Stop
Relationships	1. Start 2. Stop
BenchMark	1. Start 2. Stop
Scheduler	
DBConnectionCache	

Administering membership and security policy information

Administrators can import membership and security policy information to be used with the RBAC functionality from any authorized server. Conversely, membership and security policy information can also be exported to a file for use on an additional server or for storage.

Importing membership and security policy information

Perform the following steps to import membership or security policy information:

1. On the Context Menu, select Import Roles and Security Policy. This displays the Import dialog box, where you can specify the file path.
2. Select the file to import. If you import information when the User/Roles Management view is active, the changes will not display until you close and re-open the view.

Note: You may also import information using the `-xmsp` option using `repos_copy`. For information on using repos_copy, refer to “Using repos_copy” on page 108.

Exporting membership and security policy information

Perform the following steps to export membership or security policy information:

1. On the Context Menu, select Export Roles and Security Policy. This displays the Export dialog box, where you can specify the file path.
2. Select the destination for the file to export.

Note: You may also export information using the `-xmsp` option using `repos_copy`. For information on using repos_copy, refer to "Using repos_copy" on page 108.

Administering the RBAC password

Each user in RBAC has an associated password. When a user logs in to the server, the password is used to verify the roles assigned to the user. Occasionally, it may become necessary to change or reset the user password. Perform the following steps to reset the user password:

1. On the Context Menu, highlight the user for whom you'd like to reset the password.
2. Select Reset Password. This displays the Reset Password dialog box, with the username populated.
3. In the New Password field, enter the new password.
4. In the Confirm Password field, enter the new password again. The password is now reset.

Security Administration

As an administrator, you can monitor the use of the roles in RBAC using the security administration functionality. The InterChange Server lists active users in a table, which displays username, session ID, and the amount of time the user has spent logged onto the server.

Note: It is recommended that you refresh the user listing occasionally to retain an accurate user display. Refresh the user listing by selecting the Refresh option on the Context menu.

Viewing active users

Perform the following steps to view active users:

1. On the Context menu, select Security Administration. This opens a dialog box which displays all active users in table format.

Logging out active users

Perform the following steps to log active users off of the server:

1. To log the user out of all sessions, select the Log Out Context menu.
2. To log the user out of the selected session, select the Log Out Session Context menu.

Chapter 3. Administering problem scenarios

This chapter provides troubleshooting topics to help determine and resolve problem scenarios that an administrator may have to resolve using IBM WebSphere InterChange Server. For more information on troubleshooting server problems, see the IBM WebSphere InterChange Server Problem Determination Guide. The following topics are covered:

"Administering failed events"

"Administering run-time properties" on page 147

"Administering High-Availability (HA) systems" on page 148

"Administering the Object Request Broker" on page 150

Administering failed events

You can use two tools in IBM WebSphere InterChange Server to locate, view, and process failed events: Failed Event Manager, a browser-based tool with role-based security that allows you to work with failed events from the Web, and Flow Manager, a tool installed with the IBM WebSphere InterChange Server product. This section contains the following:

"Using Failed Event Manager"

"Using Flow Manager" on page 140

Using Failed Event Manager

Failed Event Manager allows you to view and manage failed events from the Web, and works with IBM WebSphere Application Server or with Tomcat. For information on installing Failed Event Manager to work with WebSphere Application Server or Tomcat, refer to the *WebSphere Business Integration Server Installation Guide*. Failed Event Manager works with the following versions of WebSphere Application Server, WebSphere Application Server, and Tomcat:

- WebSphere Application Server versions 5.0.2 or 5.1

  If you selected an Administrative Toolset installation of WebSphere Business Integration Server, Failed Event Manager is automatically installed and configured if WebSphere Application Server or WebSphere Application Server versions 5.0.2 or 5.1 are detected on your system.
- Tomcat versions 4.1.24 and 4.1.27.

  By default, role-based security is enabled after you have installed Tomcat with Failed Event Manager (Refer to the *WebSphere Business Integration Server Installation Guide*). Default roles must be added in `*Tomcat_home*\conf\tomcat-users.xml` directory. Creating a user with the Administrator role allows that user to gain full access to Failed Event Manager. For details on how to create roles in Failed Event Manager, refer to "Steps for creating custom users and roles for Failed Event Manager with Tomcat" on page 135

Role-based security is available if you are using Tomcat with Failed Event Manager. The actions you can perform with those failed events depends on role-based security. Administrators can assign users one or more of four default roles, and administrators have permission to create custom roles for your specific team using Application Assembly Tool. The following default roles are included:

- Administrator -- rights to view and resubmit failed events and to view business object data.
- ViewAll -- rights to view events and business object data.
- ViewEvents -- rights to view events only, not business object data.
- SubmitEvents -- rights to view and resubmit events but not business object data.

This section includes the following procedures:

- "Steps for manually installing Failed Event Manager on WebSphere Application Server 5.0.2"
- "Steps for manually installing Failed Event Manager on Tomcat 4.1.24" on page 133
- "Steps for creating custom users and roles for Failed Event Manager with Tomcat" on page 135
- "Steps for logging on to Failed Event Manager" on page 136
- "Steps for viewing failed events" on page 137
- "Steps for processing failed events in Failed Event Manager" on page 139
- "Steps for checking your access rights in Failed Event Manager" on page 140

Steps for manually installing Failed Event Manager on WebSphere Application Server 5.0.2

Perform the following steps to manually install Failed Event Manager on WebSphere Application Server 5.0.2. Prior to beginning the steps below, ensure that WebSphere Application Server is running in Administrator mode on your environment.

1. With the WebSphere Application Server running, expand the Servers menu in the left navigation frame of the Administration Console.
2. Select the application server to configure under the `Application Servers` link.
3. Select `Additional Properties > Process Definition > Java Virtual Machine`.
4. In the Generic JVM Arguments dialog, enter the following information:
   ```
       -DORBNamingProvider=CosNaming -
   Dorg.omg.CORBA.ORBClass=com.ibm.CORBA.iiop.ORB -
   Dorg.omg.CORBA.ORBInitialPort=%ORB_PORT% -
   Dorg.omg.CORBA.ORBInitialHost=%ORB_HOST%
   ```
   where `%ORB_PORT%` & `%ORB_HOST%` matches the information included in the `./bin/CWSharedEnv.bat` file of the applicable WebSphere InterChange Server.
5. Select the **Apply** button.
6. Select `Additional Properties > Custom Properties > New`.
7. In the `Name` field, enter `FEM_HOME` in the Name field to designate where the log file will be stored.
8. In the `Value` field, enter the fully-qualified path to the installed application within the WebSphere product directory, for example,
   `C:\ProgramFiles\IBM\WebSphere\Express\AppServer\installedApps\DefaultNode\FailedEvents.war.ear\FailedEvents.war.`

   Note: If the application is not installed already, either complete steps 9-20 and come back to this step or type in the fully qualified path to a directory where you would like the log file to be placed
9. Select the **Apply** button.
10. Select the **OK** button to return to the Java Virtual Machine page.

11. Select the **Save** button when the following message appears:
    `Changes have been made to your local configuration. Click Save to apply changes to the master configuration.`
12. Select the **Save** button again on the Master Configuration screen.
13. Select `Environment > Update Web Server Plugin`.
14. Select the **OK** button on the Update web server plugin configuration screen.
15. On the Applications menu, select `Install New Application`.
16. Enter the fully qualified path to the `FailedEvents.war` file located in the `WebSphereICS\WBFEM` directory.
17. Enter `FailedEvents` in the Context Root field and select the **Next** button.
18. Select the **Next** button on the `Preparing for the application installation` window to accept the default values.
19. Select the **Next** button until you reach the end of the installation process, and then select the **Finish** button.
20. When you receive the `Application Installed Successfully` message, select the **Save to Master Configuration** button.
21. Select `Start > Programs > IBM WebSphere Application Server v5.0 > Stop the Server` to stop the WebSphere Application Server. Restart the server entering the following information at a command prompt:
    ```
    WAS_Product_dir\bin\startServer.bat <servername>
    ```
22. To start the Failed Event Manager, open a browser window and enter the following URL:
    `http://hostname:9080/FailedEvents`.

    Note: 9080 is the default port for server1. To find specific ports for a server, see “Steps for identifying server ports.”

Steps for identifying server ports: Perform the following steps to identify specific server ports for use with the Failed Event Manager:

1. With the WebSphere Application Server running, expand the Servers menu in the left navigation frame.
2. Select `Application Servers`.
3. Select the server for which you require port information to expand the listing.
4. Select `Additional Properties > WebContainer`.
5. Select `Additional Properties > HTTP Transports`. The ports for the server display on the screen.

Steps for manually installing Failed Event Manager on Tomcat 4.1.24

Perform the following steps to manually install Failed Event Manager on Tomcat 4.1.24: Prior to beginning the steps below, ensure that WebSphere Application Server is running in Administrator mode on your environment.

1. Under the `Tomcat_home\webapps` directory, create the `FailedEvents` directory, where `Tomcat_home` is the Tomcat installation path.
2. Extract the `FailedEvents.war` file contents from the the `\WBFEM\Tomcat` directory into the `Tomcat_home\webapps\FailedEvents` directory.
3. Open the `setclasspath.bat` file, located in `Tomcat_home\bin` directory and set the `JAVA_OPTS` property as follows:
   ```
   set JAVA_OPTS=-DFEM_HOME=C:\Tomcat_home\webapps\FailedEvents -
   DORBNamingProvider=CosNaming -
   Dorg.omg.CORBA.ORBClass=com.ibm.CORBA.iiop.ORB -
   ```

```
Dorg.omg.CORBA.ORBInitialPort=%ORB_PORT% -
Dorg.omg.CORBA.ORBInitialHost=%ORB_HOST%
```

where `%ORB_PORT%` & `%ORB_HOST%` match what is in the `./bin/CWSharedEnv.bat` file.

Note: If `Tomcat_home` contains spaces, use quotes around the `FEM_HOME` value.

4. Start Tomcat using the following command line to ensure that `setclasspath.bat` is called:
   `Tomcat_home/bin/startup.bat`
5. With the application server running, start the Failed Event Manager by opening a browser window and entering the following URL: `http://hostname:8080/FailedEvents`.

   Note: 8080 is the default port for the server. To change the port number, you must edit the `Tomcat_home\conf\server.xml` file and restart the application server.

Role-based security

The Failed Event Manager provides the ability to activate role-based security on WebSphere Application Server and Tomcat. For more information on the concepts and application of role-based access, see "Administering role-based access control (RBAC)" on page 123.

Four basic roles exist for the Failed Events Manager.

- Administrator, which is the default role provided by the WebSphere InterChange Server and grants all privileges to the user
- SubmitEvents, which grants the user the ability to view and manage, that is, submit and delete, events. Users with this permission cannot view business object data.
- ViewAll, which grants the user the ability to only view events and business object data.
- ViewEvents, which grants the user the ability to only view events. Users with this permission cannot view business object data.

The Administrator will create all but the default Administrator role, giving each View Failed Events permission at the server level. You may also create custom roles for use with the Failed Events Manager. For more information, see "Creating custom roles for Failed Event Manager"

Note: You must use the SubmitEvents role along with another custom role that has permission to submit & delete failed events for collaborations and connectors at the component level. The ViewAll and ViewEvents roles can be used independently of another custom role.

Creating custom roles for Failed Event Manager: You can create custom roles in order to provide component level access. For example, using custom roles at a component level, you can restrict a user from viewing failed events that belong to a specific event owner, collaboration, or connector. Create custom roles using the Server Administration Tool in the System Manager., assigning View Failed Events permission at the server level. When creating roles at the component level, select the particular event owner, collaboration object or connector object and assign both Submit Failed Events and Delete Failed Events permission for each component.

Steps for creating custom users and roles for Failed Event Manager with Tomcat

Perform the following steps to create custom users and roles if you are using Tomcat:

1. Edit the `tomat-users.xml` file located under `Tomcat_Home\conf` to make the following changes as needed:
   - To add a new role, for example, *Manager*, add `<role rolename="`*`Manager`*`"/>`.
   - To create and assign a user, for example, *Scott,* to the role *Manager*, add `<user username="`*`Scott`*`" password="`*`tiger`*`" roles="`*`Manager`*`"/>`
   - You can assign more than one role to a user by separating the roles with commas, for example: `<user username="`*`Scott`*`" password="`*`tiger`*`" roles="`*`Manager, Employee`*`"/>`
2. Save the file.
3. Edit the `web.xml` file located under `Tomcat_Home\webapps\FailedEvents\WEB-INF` to add the roles you added to the `tomat-users.xml` file in step 1. Do the following:
   - Add the custom roles to the following xml element, which already contains the default roles:

     ```
     <auth-constraint id="AuthConstraint_1062537631424">
                 <description>SC1:+:</description>
                 <role-name>Administrator</role-name>
                 <role-name>ViewEvents</role-name>
                 <role-name>ViewAll</role-name>
                 <role-name>SubmitEvents</role-name>
                 <role-name>Manager</role-name>
              </auth-constraint>
     ```

   - At the end of the xml file add a new element and assign a security role id that is unique:

     ```
     <security-role id="SecurityRole_1068513225089">
              <description>Can manage all events.</description>
              <role-name>Manager</role-name>
      </security-role>
     ```

   - Edit the init param value of the Login servlet as one complete string:

     ```
     <servlet id="Servlet_1062537018298">
     <servlet-name>Login</servlet-name>
     <display-name>Login</display-name>

     <servlet-class>
     com.ibm.btools.itools.FailedEvents.servlets.Login
     </servlet-class>
     <init-param id="InitParam_1063835207426">
     <param-name>ROLECOMPONENTS</param-name>
     <param-value>
     role1: event_owners
     SourceToDestCollab
     |role2:event_owners=Collab2*
     bos=CUSTOMER2.Create#2/4
     |role3:event_owners=SourceToDestCollab, Collab2
     </param-value>
     <description>
     Roles and components (collabs or connectors) associated with these roles.
     </description>
     </init-param>
     </servlet>
     ```

     The parameter value is a string in the following format:

     ```
     <RoleName1>:event_owners=<ownername>*connectors=<connectorName>
     *bos=<boname.verb> #<CompositeKeyValue>+<CompositekeyValue2>/
     <AnotherPossibleKeyValue> | <RoleName2>: .....
     ```

Where:

< | > – separates two roles

< : > – separates the role name from components of the role

< * > – separates components within a role, for example, connectors and business objects

< = > - separates the component name from its values

< , > - separates values within a component

< # >- separates business object name from its key values

< / >- separates different possible keys of a business object

< + >- separates the composite key of a business object (two or more primary keys)

For example:

```
Role1:event_owners=collab1, collab2*connectors=conn1,
conn2*bos=bo1.create#55/67, bo2.delete#99/80 |
```

```
Role2:event_owners=collab3,collab4*connectors=conn3
*bos= bo4.create#59+9876/82, bo2.delete#56
```

In this example, “event_owners” refers to point of failure of the event, either at the collaboration or at the connector.

4. Save the file.
5. Start Tomcat from the command line in order to call the setclasspath.bat file: `Tomcat_home/bin/startup.bat`.

Steps for logging on to Failed Event Manager

Perform the following steps to log on to Failed Event Manager:

1. Type one of the following URLs, depending on how Failed Event Manager was set up for your team:
   - If you are using WebSphere Application Server, type the URL:

     ```
     http://HostName/FailedEvents
     ```

     where *HostName* is the name of the computer on which WebSphere Application Server is installed.
   - You may have a port number as part of the URL. To access Failed Event Manager, type the URL:

     ```
     http://HostName:nnnn//FailedEvents
     ```

     where *HostName* is the name of the computer on which WebSphere Application Server is installed and *nnnn* is the port number.
   - When Failed Event Manager is installed as part of WebSphere Application Server, it is configured to use a default port number, 7089. You can access Failed Event Manager with the following URL: `http://HostName:7089/FailedEvents`.
   - If you are using Failed Event Manager with Tomcat, it is configured to use a default port number, 8080. You can access Failed Event Manager with the following URL: `http://HostName:8080/FailedEvents`.
2. If security is enabled, you must type the Application Server user name and password.
3. At the Connect screen, type the **Server Name**, **User Name** and **Password**, and click **Login**.

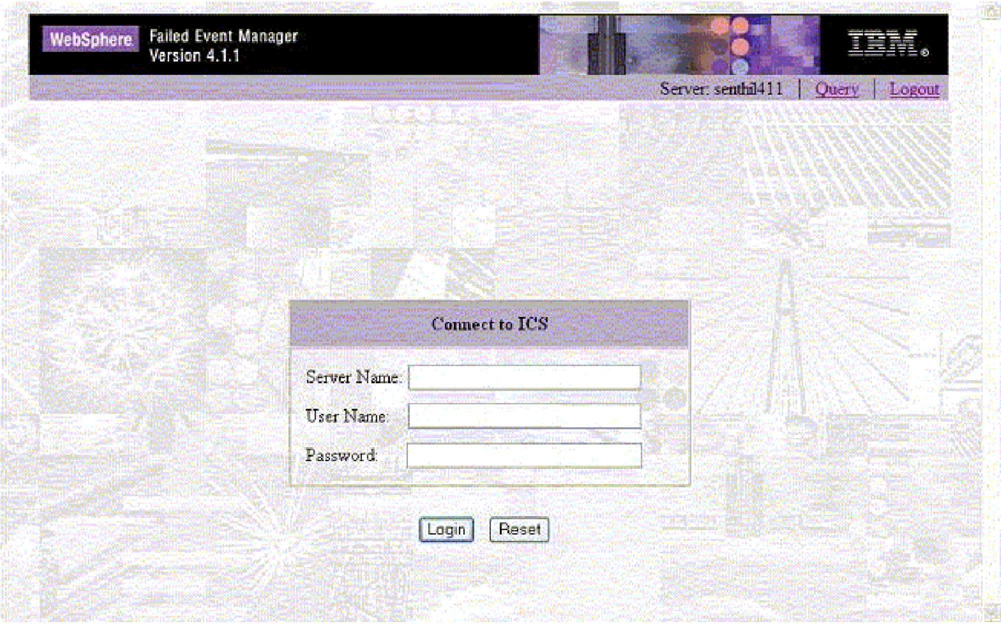

Figure 45. Failed Event Manager Connect screen

After you log on, the Query dialog box appears, as shown in Figure 46.

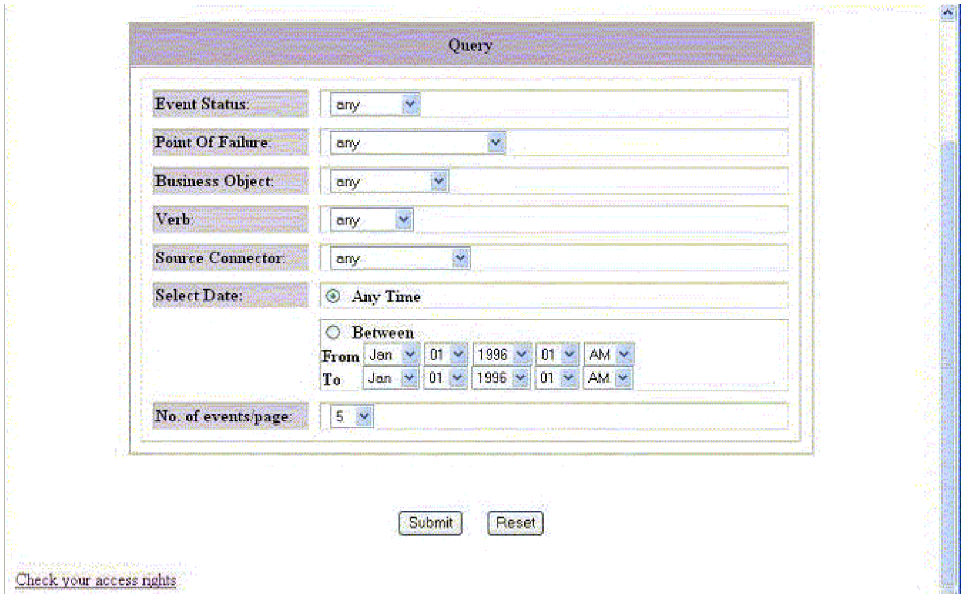

Event Status: any

Point Of Failure: any

Business Object: any

Verb: any

Source Connector: any

Select Date: ● Any Time

○ Between

From Jan 01 1996 01 AM

To Jan 01 1996 01 AM

No. of events/page: 5

Figure 46. Failed Event Manager Query dialog box

Steps for viewing failed events

Perform the following steps to view information about failed events in Failed Event Manager:

1. On the Query page, select the information about the failed events that you want to view:
   - **Event Status**: select whether you want to view events with the status of **Failed**, **In Transit**, **Possible Duplicate**, **Deferred**, or **Any** of those four status categories.
   - **Point of Failure**: select the component in which the failure occurred.

- **Business Object**: select a name of a business object.
- **Source Connector**: select a name of an adapter or other source connector.
- **Select Date**: click **Any Time**, or click Between and select a beginning and end date and time.
- **No. of events/page**: Select how may events you want to be displayed on each page of Failed Event Manager.

2. Click **Submit**.

   A Table of Failed Events appears, as show in Figure 47.

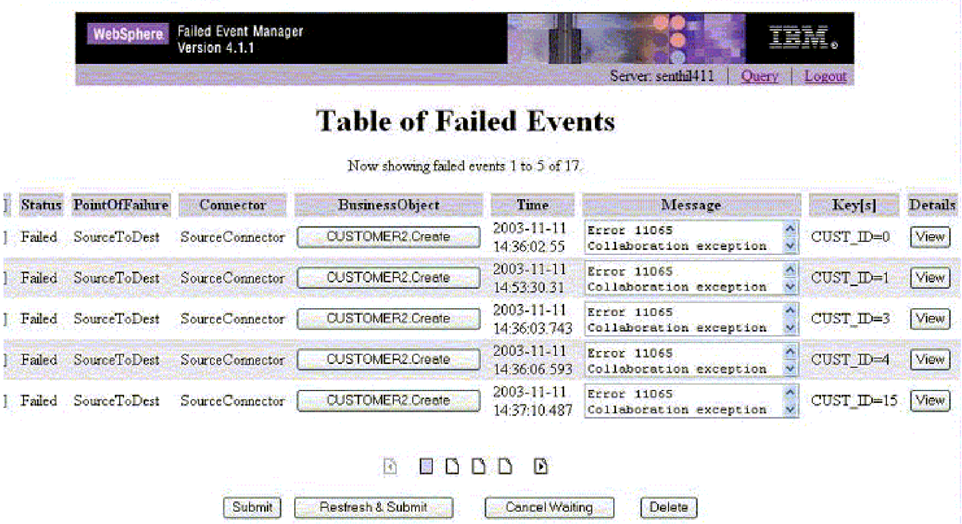

Table of Failed Events

Now showing failed events 1 to 5 of 17.

Status	PointOfFailure	Connector	BusinessObject	Time	Message	Key[s]	Details
Failed	SourceToDest	SourceConnector	CUSTOMER2.Create	2003-11-11 14:36:02.55	Error 11065 Collaboration exception	CUST_ID=0	View
Failed	SourceToDest	SourceConnector	CUSTOMER2.Create	2003-11-11 14:53:30.31	Error 11065 Collaboration exception	CUST_ID=1	View
Failed	SourceToDest	SourceConnector	CUSTOMER2.Create	2003-11-11 14:36:03.743	Error 11065 Collaboration exception	CUST_ID=3	View
Failed	SourceToDest	SourceConnector	CUSTOMER2.Create	2003-11-11 14:36:06.593	Error 11065 Collaboration exception	CUST_ID=4	View
Failed	SourceToDest	SourceConnector	CUSTOMER2.Create	2003-11-11 14:37:10.487	Error 11065 Collaboration exception	CUST_ID=15	View

Figure 47. Failed Event Manager Table of Failed Events

If you are not assigned to one of the roles that has access to viewing business objects, the business object buttons may not be available in this table, as shown in Figure 48.

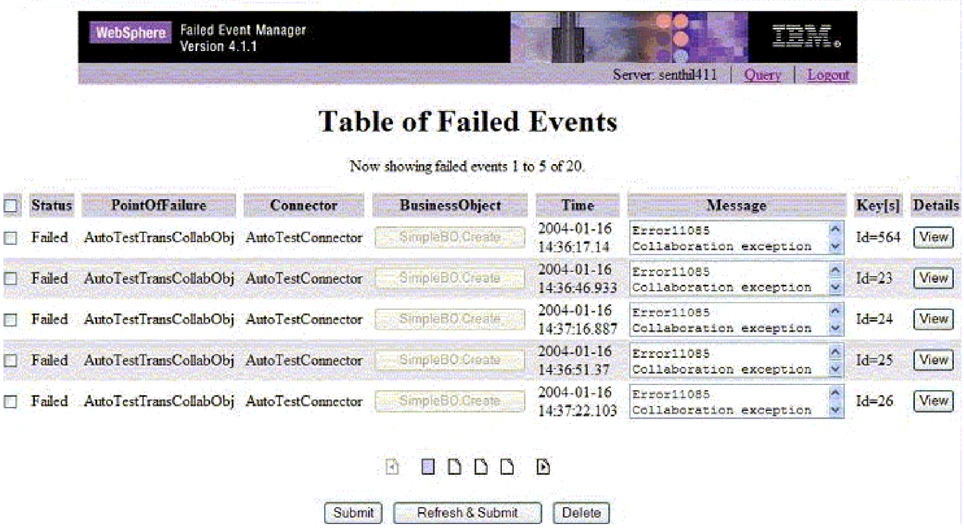

Table of Failed Events

Now showing failed events 1 to 5 of 20.

□	Status	PointOfFailure	Connector	BusinessObject	Time	Message	Key[s]	Details
□	Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:36:17.14	Error11085 Collaboration exception	Id=564	View
□	Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:36:46.933	Error11085 Collaboration exception	Id=23	View
□	Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:37:16.887	Error11085 Collaboration exception	Id=24	View
□	Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:36:51.37	Error11085 Collaboration exception	Id=25	View
□	Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:37:22.103	Error11085 Collaboration exception	Id=26	View

Figure 48. Failed Event Manager Table of Failed Events without rights to view business objects

3. If you want to view details for a particular failed event, select the check box in the left column and click **View** in the **Details** column.

The Event Details dialog box appears, as shown in Figure 49.

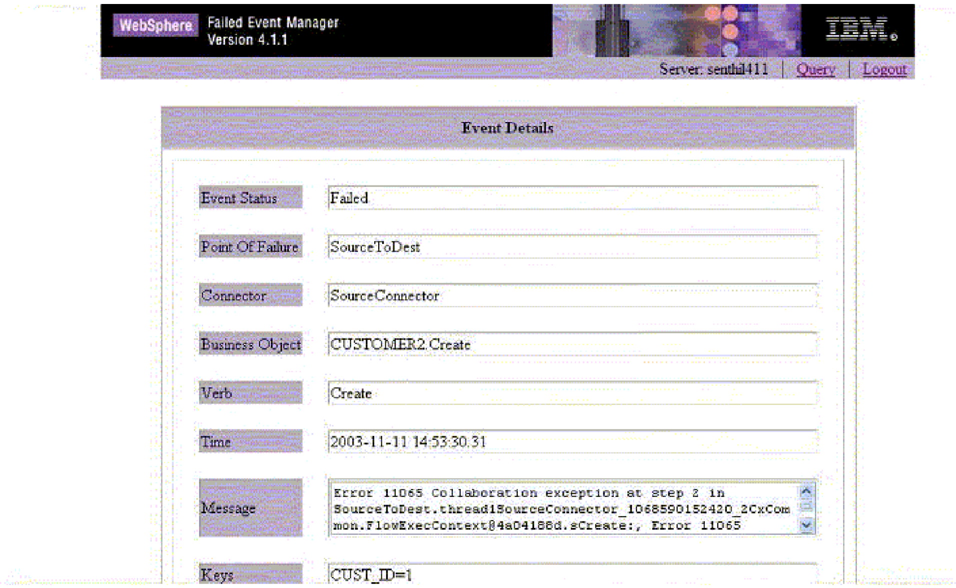

Event Details

Event Status	Failed
Point Of Failure	SourceToDest
Connector	SourceConnector
Business Object	CUSTOMER2.Create
Verb	Create
Time	2003-11-11 14:53:30.31
Message	Error 11065 Collaboration exception at step 2 in SourceToDest.thread1SourceConnector_1068590152420_2CxCommon.FlowExecContext@4a04188d.sCreate:, Error 11065
Keys	CUST_ID=1

Figure 49. Failed Event Manager Event Details dialog box

4. If you are assigned to a role that has access to the business object, click the business object button in the **Business Objects** column to view details about the business object. The Business Object Data dialog box appears, as shown in Figure 50.

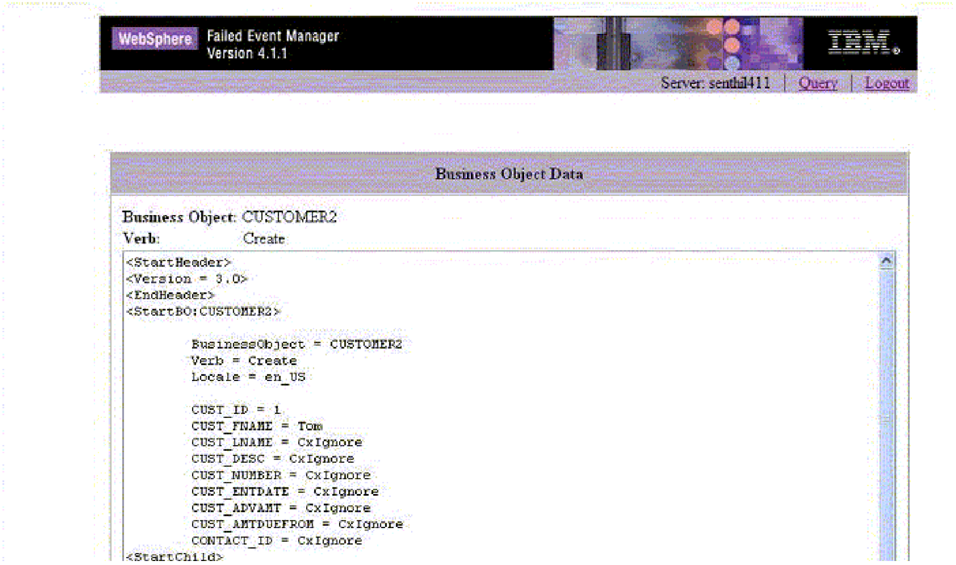

Business Object Data

Business Object: CUSTOMER2
Verb: Create

```
<StartHeader>
<Version = 3.0>
<EndHeader>
<StartBO:CUSTOMER2>

        BusinessObject = CUSTOMER2
        Verb = Create
        Locale = en_US

        CUST_ID = 1
        CUST_FNAME = Tom
        CUST_LNAME = CxIgnore
        CUST_DESC = CxIgnore
        CUST_NUMBER = CxIgnore
        CUST_ENTDATE = CxIgnore
        CUST_ADVAMT = CxIgnore
        CUST_AMTDUEFROM = CxIgnore
        CONTACT_ID = CxIgnore
<StartChild>
```

Figure 50. Failed Event Manager Business Object Data dialog box

Steps for processing failed events in Failed Event Manager

Perform the following steps in a Table of Failed Events to process failed events:

1. Select the check box in the left column for the failed events that you want to work with.
2. Do one of the following:

- To send the events back to the destination application, click **Submit**. This action is available if you are assigned to a role that can submit events.
- To refresh the selected events and then send them back to the destination application, click **Refresh & Submit**.This action is available if you are assigned to a role that can submit events.
- To cancel the selected events if they are long-lived business processes, click **Cancel Waiting**.
- To delete the selected events, click **Delete**.

Steps for checking your access rights in Failed Event Manager

Perform the following steps to check the roles you are assigned for Failed Event Manager role-base security:

1. log on to Failed Event Manager (see “Steps for logging on to Failed Event Manager” on page 136).
2. On the query page that opens, click the **Check your access rights** link in the bottom left.

   A page appears listing the default roles available. You can type a role name to see if you are assigned to that role.

Using Flow Manager

To locate, view, and process failed events, use Flow Manager, a tool that is installed with the IBM WebSphere InterChange Server product. Flow Manager allows you to easily construct a query to locate and display unresolved flows. After you display the unresolved flows, you can select any flow in the display and submit it, discard it, or perform other actions.

The following topics describe how to use Flow Manager for constructing the queries, viewing the details, and processing the events:

“Steps for starting Flow Manager”

“Steps for finding unresolved flows” on page 142

“Steps for viewing details for unresolved flows” on page 145

“Steps for processing unresolved flows” on page 147

Steps for starting Flow Manager

Perform the following steps to start Flow Manager:

1. Click **Start > Programs > IBM WebSphere Business Integration > Toolset > Administrative > Flow Manager**. The Connect to WebSphere InterChange Server dialog box appears (see Figure 51 on page 141.

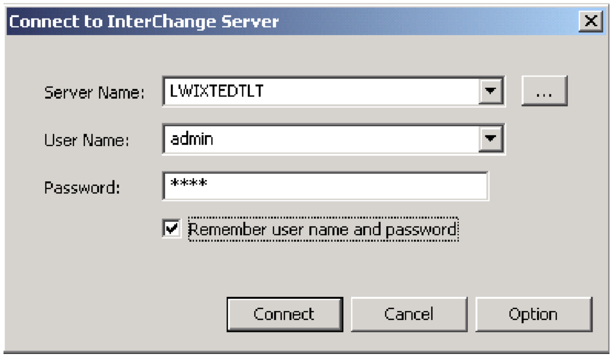

Server Name: LWIXTEDTLT

User Name: admin

Password: ****

☑ Remember user name and password

Figure 51. Connect to InterChange Server dialog box

2. Enter the server name, using any one of the following methods:
   - Type the name of the server in the **Server Name** field.
   - Select the server name from the **Server Name** list.
   - Click the browse button to browse for the server on the network. The following figure shows the Server dialog box that opens when you click this button.

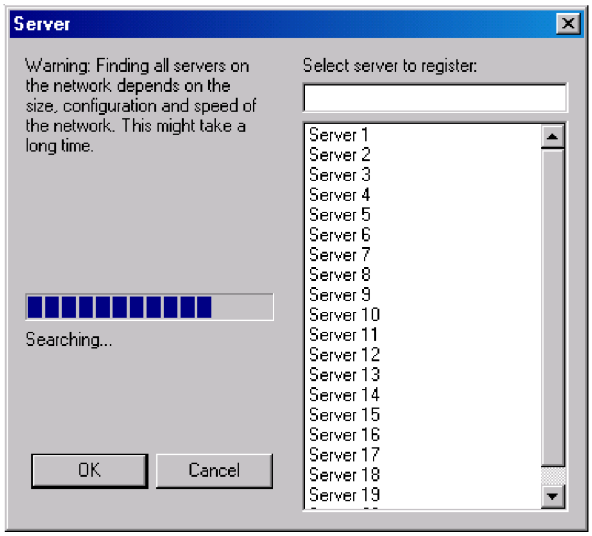

Warning: Finding all servers on the network depends on the size, configuration and speed of the network. This might take a long time.

Searching...

Server 1
Server 2
Server 3
Server 4
Server 5
Server 6
Server 7
Server 8
Server 9
Server 10
Server 11
Server 12
Server 13
Server 14
Server 15
Server 16
Server 17
Server 18
Server 19

Figure 52. Server dialog box

Note: Browsing for the server on the network could take a long time, depending on how many servers exist on the network.

3. Type the user name and password for the server you want to connect to. If you check the **Remember user name and password** check box, the user name and password are stored in the registry along with the server name, and the password are encrypted. Click **Connect**.

   Note: To remove unwanted users or servers from the registry, click Options. This opens the Options dialog box, from which you can remove servers or users.

   The following figure shows the Options dialog box.

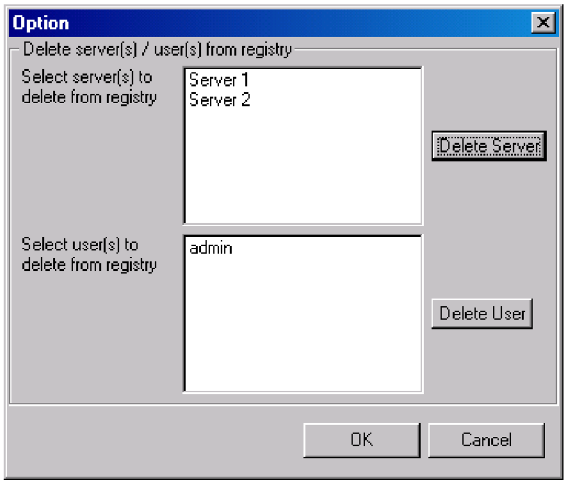

Select server(s) to delete from registry

Server 1
Server 2

Select user(s) to delete from registry

admin

Figure 53. Options dialog box

After you click **Connect** in the Connect to WebSphere InterChange Server dialog box, Flow Manager opens (see Figure 54).

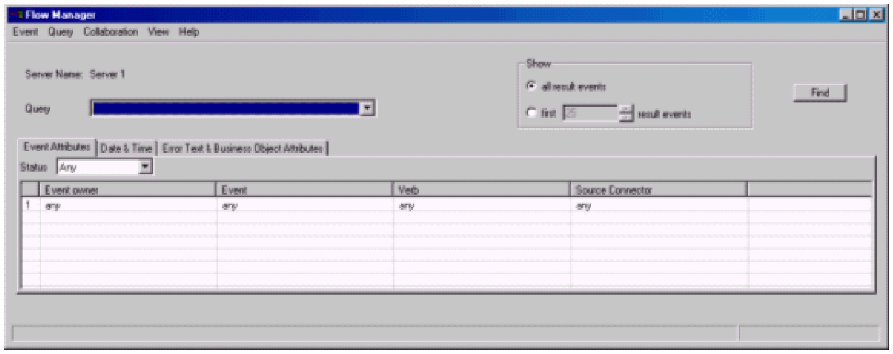
Flow Manager
Event Query Collaboration View Help
Show
Find
Event Attributes
Date & Time
Error Text & Business Object Attributes

Server Name: Server 1

Query

● all result events

○ first 25 result events

Status: Any

	Event owner	Event	Verb	Source Connector
1	any	any	any	any

Figure 54. Flow Manager

Steps for finding unresolved flows

Perform the following steps in Flow Manager to find all unresolved flows:

1. Start Flow Manager. See “Steps for starting Flow Manager” on page 140.
2. Do one of the following:
   - If you want to locate and display all unresolved flows, accept the defaults of **Any** for all the fields.
   - If you want to construct a query to find specific flows, refer to the following sections that describe the search options in each tab:
     - “Event Attributes tab” on page 143
     - “Date & Time tab” on page 144
     - “Error Text & Business Object Attributes tab” on page 144
3. After specifying your search criteria, click the **Find** button or click **Event > Find**. The results appear in the bottom half of the Flow Manager window (see Figure 55 on page 143).

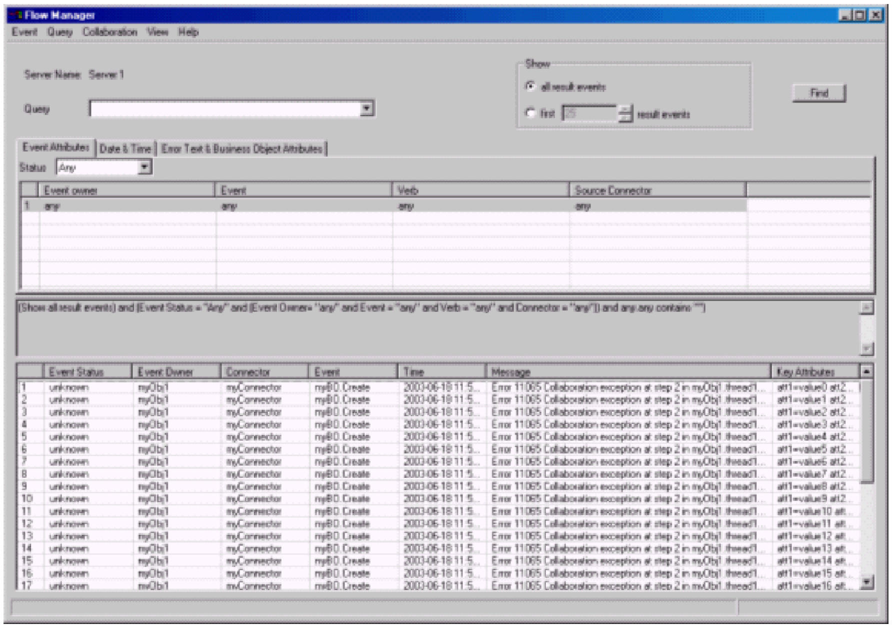
Flow Manager
Event Query Collaboration View Help
Show
Find
Event Attributes | Date & Time | Error Text & Business Object Attributes

Server Name: Server 1

Query:

● all result events
○ first 25 result events

Status: Any

	Event owner	Event	Verb	Source Connector
1	any	any	any	any

(Show all result events) and (Event Status = "Any" and (Event Owner= "any" and Event = "any" and Verb = "any" and Connector = "any")) and any.any contains "")

	Event Status	Event Owner	Connector	Event	Time	Message	Key Attributes
1	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value0 att2...
2	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value1 att2...
3	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value2 att2...
4	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value3 att2...
5	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value4 att2...
6	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value5 att2...
7	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value6 att2...
8	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value7 att2...
9	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value8 att2...
10	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value9 att2...
11	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value10 att...
12	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value11 att...
13	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value12 att...
14	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value13 att...
15	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value14 att...
16	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value15 att...
17	unknown	myObj1	myConnector	myBO.Create	2003-06-18 11:5...	Error 11065 Collaboration exception at step 2 in myObj1.thread1...	att1=value16 att...

Figure 55. Flow Manager displaying filtered results

4. To save the current filter information as a query, click **Query > Save**. The query name you type appears in the **Query** list.

Event Attributes tab: Click the **Event Attributes** tab to search according to the characteristics of the unresolved event (see Figure 56).

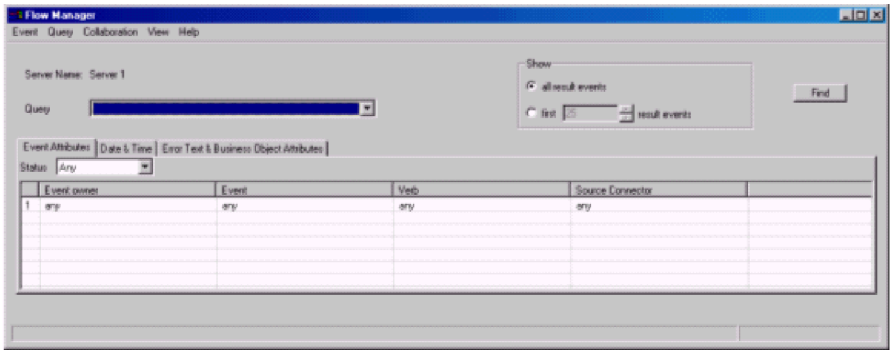
Flow Manager
Event Query Collaboration View Help
Show
Find
Event Attributes | Date & Time | Error Text & Business Object Attributes

Server Name: Server 1

Query:

● all result events
○ first 25 result events

Status: Any

	Event owner	Event	Verb	Source Connector
1	any	any	any	any

Figure 56. Flow Manager Event Attributes tab

The following options are available from the **Event Attributes** tab:

- **Status**: Select events according to their status, which you can designate as one of the following:
  - **Any**: Selects all unresolved events.
  - **Failed**: Selects all unresolved events that are in the Failed state.
  - **In Transit**: Selects all unresolved events that are in the In Transit state.
  - **Possible Duplicate**: Selects all unresolved events that are in the Possible Duplicate state.
  - **Deferred**: Selects all unresolved events in which the point of failure was a collaboration with a recovery mode setting of Deferred.
  - **Waiting**: Selects all unresolved events in which the point of failure was a collaboration with a recovery mode setting of Waiting.

- **Event Owner**: Select a collaboration name or a connector name to query for flows that failed within that collaboration or connector. You can string together multiple selections by clicking in the new empty field that is created below each existing selection. To search all of your collaborations and connectors, select **Any**.
- **Event**: For each collaboration or connector that you select, select an event, as represented by a business object. You can select one event, string together multiple events, or select **Any**.
- **Verb**:

  Chose the verb to query for each event.
- **Source Connector**

  Select the source connector or connector whose flows you want to query, or select **Any**.

Date & Time tab: Use the **Date & Time** tab to query only the flows that failed during a specific time period (see Figure 57).

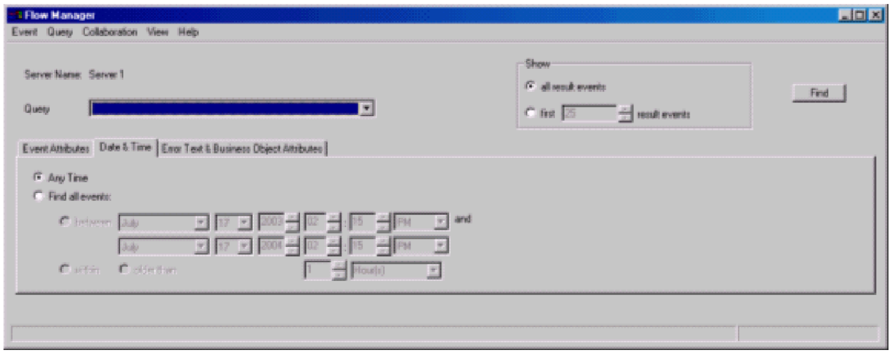
Flow Manager
Event Query Collaboration View Help
Find
Event Attributes
Date & Time
Error Text & Business Object Attributes

Server Name: Server 1

Query

Show
● all result events
○ first 25 result events

● Any Time
○ Find all events:
○ between July 17 2003 02 : 15 PM and
July 17 2004 02 : 15 PM
○ within ○ older than 1 Hour(s)

Figure 57. Date & Time tab

The following options are available from the **Date & Time** tab:

- **Any Time**: Queries for all failed events, with no restriction on when occurred
- **Find all events**: Queries for all failed events that occurred within a specified time period:
  - **between** (specified times)
  - **within** (period of days)
  - **older than** (number of days)

Error Text & Business Object Attributes tab: Use the **Error Text & Business Object Attributes** tab to construct a query that includes only certain business objects, business object attributes, and attribute values (see Figure 58 on page 145).

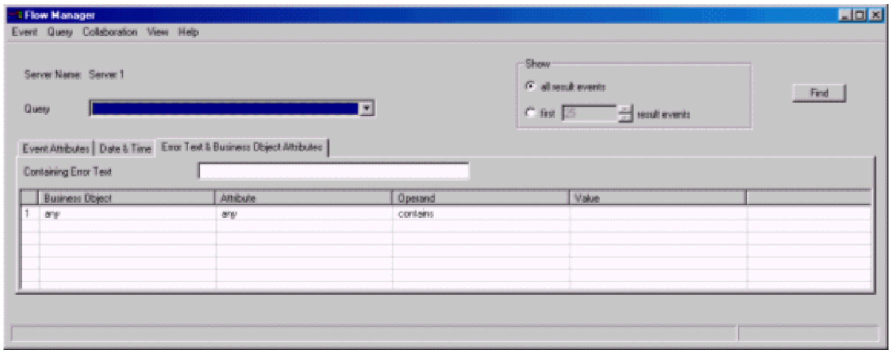
Flow Manager
Event Query Collaboration View Help
Query
Find
Event Attributes
Date & Time
Error Text & Business Object Attributes
Containing Error Text

Server Name: Server 1

Show
● all result events
○ first 25 result events

	Business Object	Attribute	Operand	Value	
1	any	any	contains		

Figure 58. Error Text & Business Object Attributes tab

Select values in the following columns:

- **Business Object**
- **Attribute**
- **Operand**
- **Value**

You can also query only unresolved flows that produce a message containing specified text that you type in the **Containing Error Text** field.

Note: The keywords you type in the **Containing Error Text** field are not case-sensitive.

Steps for viewing details for unresolved flows

Perform the following steps to view details for unresolved flows in Flow Manager:

1. Select a query from the **Query** list.
2. Do one of the following:
   - Click the **Find** button.
   - Click **Event > Find**.

   The results appear in the bottom half of the Flow Manager window (see Figure 55 on page 143).

   An unresolved flows table appears with a list of events and the following information:

   - **Event Status**
   - **Event Owner**
   - **Point of Failure**
   - **Connector**
   - **Event**
   - **Time**
   - **Message**
   - **Key Attributes**

   For the waiting events, the **Event Status** cell contains the following information:

   ```
   Waiting
   [timeout expiration:...]
   [scenario name:...]
   [scenario node ID...]
   ```

3. To access more information about any of the events, double-click the row containing the event (or click **Event > Select All**). This opens the Show Event Details dialog box (see Figure 59).

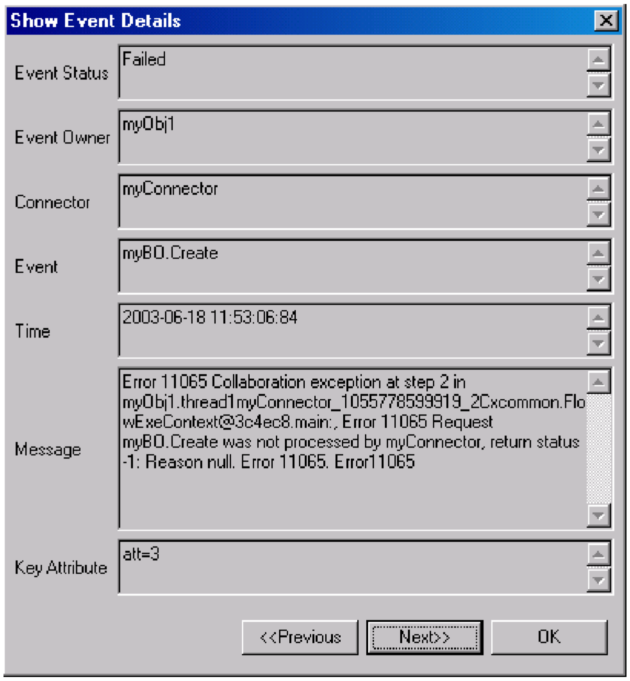

Event Status	Failed
Event Owner	myObj1
Connector	myConnector
Event	myBO.Create
Time	2003-06-18 11:53:06:84
Message	Error 11065 Collaboration exception at step 2 in myObj1.thread1myConnector_1055778599919_2Cxcommon.FlowExeContext@3c4ec8.main:, Error 11065 Request myBO.Create was not processed by myConnector, return status -1: Reason null. Error 11065. Error11065
Key Attribute	att=3

Figure 59. Flow Manager Show Event Details dialog box

4. To access more information about the business object associated with the selected event, click **Event > Display Details**. The Show Business Object Data dialog box appears (see Figure 60).

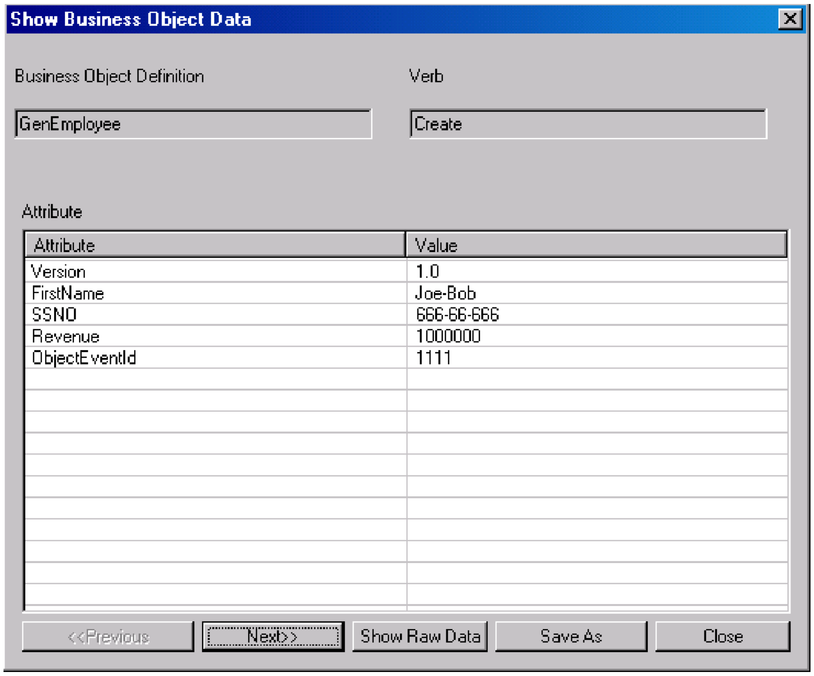

Business Object Definition	GenEmployee
Verb	Create

Attribute

Attribute	Value
Version	1.0
FirstName	Joe-Bob
SSNO	666-66-666
Revenue	1000000
ObjectEventId	1111

Figure 60. Show Business Object Data dialog box

If you want to take action on any of the unresolved flows, refer to “Steps for processing unresolved flows.”

Steps for managing queries

Perform one of the following steps to manage queries you created for unresolved flows:

- To save the current filter information as a query, click **Query > Save**. The query you type appears in the **Query** list.
- To delete a selected query in the **Query** menu, click **Query > Delete**.
- To show or hide the content of the query, click **Query > Show/Hide**.

Steps for managing results of queries

Perform the following steps to manage queries you created for unresolved flows:

1. In the results list in Flow Manager, select one or more events in the results list.
2. Do one of the following:
   - To delete selected events, click **Event > Delete**.
   - To save all the failed events in the results list in an Excel file, click **Event > Save**.
   - To print the selected events to a table, click **Event > Print**. A dialog box appears with **Print All Events** or **Print Selected Events** options.
   - To display the number of events corresponding to the current query, click **Event > Get Count**.
   - To clear all the result events, close the result list, and refresh all the filter controls, click **Event > New Search**.

Steps for processing unresolved flows

Perform the following steps to process any failed events in the results list in Flow Manager:

1. In the results list in Flow Manager, select one or more events in the results list.
2. Do one of the following:
   - To send the data of the events back to the destination application, click **Event > Submit**.
   - To cancel waiting events (long-lived business processes), click **Event > Cancel Waiting**.
   - To refresh the data and then send the events back to the destination application, click **Event > Refresh & Submit**.

Administering run-time properties

The components of the IBM WebSphere InterChange Server system obtain much of their initialization information from a single global environment file: `*ProducDir*\bin\CWSharedEnv.bat`.

The startup scripts of the various components in the IBM WebSphere InterChange Server system read this file as part of their initialization process:

- An adapter
- IBM WebSphere InterChange Server
- repos_copy utility
- Many IBM WebSphere InterChange Server tools

This CWSharedEnv file contains initialization information for the communication and run-time software that the IBM WebSphere InterChange Server system uses. Most of this information is provided in the CWSharedEnv file that comes with the product and never needs to change. Some information is customized for your system during installation.

Steps for changing run-time properties after installation

Perform one of the following steps to change run-time information *after* installation:

- Edit the appropriate variable in the CWSharedEnv file. When you change the property in the CWSharedInv file, you change it for *all* components that read this file during their setup process.
- Specify the property and its value as a command-line option to the component's startup script. When you change the property on the command line, you change it only for the component that you are starting. Properties you specify on the command line override any other property settings within the system or from the CWSharedEnv file.

Note: For more information on how to change the run-time properties that the IBM Java ORB supports, see "Steps for customizing the Object Request Broker" on page 151.

Administering High-Availability (HA) systems

A high-availability (HA) system is made up of two or more machines (the primary and one or more backup machines) that are configured identically and designated as a cluster. Each machine is considered a node in the cluster. The primary and backup nodes share a cluster name and IP address. External processes use this name and IP address to access a service on the cluster, which runs on either the primary or one of the backup nodes. All nodes have access to a shared Redundant Array of Independent Disks (RAID) storage system. For windows systems only, the shared RAID storage system is used only by the active node.

The HA configuration provides shutdown and automatic restart of unresponsive (failed) software programs, and migration to the cluster backup node when failures on the active node are detected. The cluster backup node assumes the cluster name and IP address and automatically takes over system processing until such time as the failure is corrected on the primary node and a failback is initiated (that is, manually return processing to the original system).

This section provides the following information about how to manage a high-availability system that includes IBM WebSphere InterChange Server:

"Supported HA environments"

"Maintaining a Windows HA system" on page 149

Supported HA environments

The high-availability (HA) option is available on the following operating systems:

- UNIX-based systems -- Sample scripts with a README file are provided as part of an IBM SupportPac.

- Windows 2000 -- *The System Installation Guide for Windows* provides a set of instructions on how to set up HA using Microsoft Cluster Server (MSCS) software. An IBM Category 2 SupportPac provides some dynamic link library (.dll) files for use with MSCS.

  Note: Scripts and files that assist with the HA implementation are available as unsupported in Category 2 SupportPac. You can find the information about these IBM SupportPacs at:

http://www.ibm.com/software/integration/supportpace/
category.html#cat2

In addition, both the *System Installation Guide for UNIX* and *for Windows* provide basic instructions on how to configure the hardware and software for use in an HA environment.

Maintaining a Windows HA system

Once an HA system is set up according to the configuration instructions provided in the *System Installation Guide for Windows*, it should need minimal maintenance or reconfiguration. This section summarizes some of the tasks for maintaining an HA system that is set up on a Windows operating system to use the Microsoft Cluster Server (MSCS) software. The following topics are covered:

- "Checking cluster status"
- "Detecting a failover"
- "Moving groups to perform maintenance" on page 150
- "Changing the status of a resource" on page 150

Checking cluster status

The MSCS Cluster Administrator is the primary administration tool that you use to administer and check the status of the cluster. Each resource, such as IBM WebSphere InterChange Server, is listed along with its state (online or offline, failed, or online or offline pending), owner (node 1 or node 2), and the type of resource (a description such as IBM WebSphere InterChange Server, disk resource, or connector). From this window, you can see the status of the individual services and which of the cluster nodes is active.

Other Windows administrative tools provide information about the status of the cluster. In particular, check the MSCS online help and documentation for details about using the following tools to monitor the cluster:

Windows Event Viewer	View and manage System, Security, and Application event logs
Windows Services option in the Control panel	Verify that the Cluster Service is running

Detecting a failover

Several types of icons appear in MSCS Administrator beside the various listings of groups, resources, and other elements. The most important icon to recognize is the node-down icon, which indicates that a failure has occurred on a cluster node and that its groups and resources have been transferred to the surviving cluster node. A node-down icon displays in the MSCS Administrator as a cluster node icon with a red x through it.

The node-down icon does not necessarily mean that you have lost functionality in any of your groups or resources. In normal operation, the group fails over to the backup cluster node.

Moving groups to perform maintenance

When you stop the Cluster Service on a node, you prevent clients from accessing cluster resources through that node and all groups move to the other node (if the failover policies allow it). This can be useful when you need to take the primary node offline to perform maintenance or upgrade its software. The following steps describe how to perform a move to the backup node for maintaining or upgrading the primary node:

1. Stop the Cluster Service on the backup node by clicking its node icon, clicking the File menu, and selecting Stop Cluster Service.
2. Upgrade the backup node if a software upgrade is the maintenance being performed. Be sure to place executables and libraries on the node disk and data files on the shared RAID.
3. Restart the Cluster Service on the backup node by clicking File > Start Cluster Service.
4. On the primary node, right-click each group and select Move Group.

   This allows you to take a server offline without losing availability of your resources.
5. Stop the Cluster Service on the primary node by clicking its node icon, then clicking File > Stop Cluster Service.
6. Upgrade the primary node in the same manner as you upgraded the backup node.
7. Restart the Cluster Service on the primary node by clicking File > Start Cluster Service.
8. On the backup node, right-click each group and select Move Group to move the groups back to the primary node.

Changing the status of a resource

Use Cluster Administrator to manually bring individual resources online or take them offline. Change the status by selecting the resource, and from the File menu, selecting either Bring Online, Take Offline, or Initiate Failure. You should only off or online WebSphere MQ, connectors, and IBM WebSphere InterChange Server.

Administering the Object Request Broker

This section contains information about resolving problems related to the IBM Java Object Request Broker (ORB), which handles communication between IBM WebSphere InterChange Server and several of its components. This section provides information on the following topics:

- "Object Request Broker installation"
- "Steps for customizing the Object Request Broker" on page 151
- "Steps for changing the location of the Object Request Broker" on page 152
- "Using the IBM Transient Name Server" on page 153

Object Request Broker installation

IBM WebSphere InterChange Server requires the IBM Java ORB to communicate with several of its components, including adapters and System Manager. Use of the ORB requires installation of the IBM Java ORB.

The IBM Java ORB is installed as part of the IBM Java Runtime Environment (JRE) software, which the IBM WebSphere InterChange Server Installer installs automatically.

Steps for customizing the Object Request Broker

Perform the following steps to customize the Object Request Broker:

1. Refer to Table 24 on page 152 for the properties that the IBM Java ORB supports to customize its behavior.
2. Specify the IBM ORB property and its value as a command-line option to the component's startup script. When you specify the IBM ORB property on the command line, you change it only for the component that you are starting. You specify an ORB property by preceding it with the `-D` command-line option. Properties you specify on the command line override any other property settings within the system or from the `CWSharedEnv` file.
3. Edit the appropriate variable in the `CWSharedEnv` file. When you change the variable in the `CWSharedEnv` file, you change it for *all* components that read this file during their startup process. These components include any adapters, IBM WebSphere InterChange Server instances, the repos_copy utility, and the IBM WebSphere InterChange Server tools.

   Note: For more information on the `CWSharedEnv` file, see "Administering run-time properties" on page 147.

   As Table 24 on page 152 shows, the ORB location is specified by special variables in the `CWSharedEnv` file. You must modify these variables in the `CWSharedEnv` file to change its location.

   Other ORB properties are listed in the `ORB_PROPERTY` variable of the `CWSharedEnv` file. In this variable, each IBM ORB property is preceded by the `-D` command-line option. To add or change an ORB property, you must add or change the appropriate `-D` option in the `ORB_PROPERTY` variable of the `CWSharedEnv` file. Properties you specify in the `CWSharedEnv` file override any other settings from the configuration file.
4. Specify the configuration parameter (if one exists) in the appropriate configuration (.cfg) file. You can set many of the ORB properties with configuration parameters in the CORBA section of the configuration file. Both the IBM WebSphere InterChange Server configuration file (InterchangeSystem.cfg) and the adapter local configuration file can contain a CORBA section. When you specify the configuration parameter in the CORBA configuration file, you change it for all ORB-related tasks that the ORB server performs.

   Important: The configuration files are in XML format. Do *not* modify these files unless you use an XML editor or are very familiar with XML format!

   For example, to specify the maximum number of threads, you can take any of the following actions:

   - Add the IBM ORB property to the `ORB_PROPERTY` variable in the `CWSharedEnv` file:

     ```
     ORB_PROPERTY=-DORBNamingProvider=CosNaming
     -Dorg.omg.CORBA.ORBClass=com.ibm.CORBA.iiop.ORB
     -Dorg.omg.CORBA.ORBInitialPort=%ORB_PORT%
     -Dorg.omg.CORBA.ORBInitialHost=%ORB_HOST%
     -Dcom.ibm.CORBA.Debug.Output=nul
     ```

```
-Dcom.ibm.CORBA.ThreadPool.MaximumSize=100
```

- Specify the IBM ORB property on the command line when you start up the component:

  start_server.....-**Dcom.ibm.CORBA.ThreadPool.MaximumSize=100**
- Add the `OAthreadMax` configuration parameter to the CORBA section of the configuration file:

  [CORBA]

  OAthreadMax=100

Table 24. IBM Java ORB properties that can be customized

IBM ORB property	Configuration parameter	Description
`com.ibm.CORBA.ListenerPort`	`OAport`	Port number on which the ORB server (within IBM WebSphere InterChange Server) listens for incoming requests.
`com.ibm.CORBA.LocalHost`	`OAipAddr`	IP address or host name of the machine on which the ORB server (within IBM WebSphere InterChange Server) is running.
`com.ibm.CORBA.ThreadPool.` `MaximumSize`	`OAthreadMax`	Maximum number of threads that the connection manager can create. The default value (zero) indicates that no size restriction exists.
`com.ibm.CORBA.ThreadPool.` `InactivityTimeout`	`OAthreadMaxIdle`	The time (in seconds) before an idle thread is destroyed.
`com.ibm.CORBA.RequestTimeout`	*None*	Number of seconds that a CORBA request waits before timing out. By default, there is no timeout; the ORB waits indefinitely for a response.
`com.ibm.CORBA.LocateRequest`	*None*	Timeout value (in seconds) for Locate Requests.
`com.ibm.CORBA.FragmentTimeout`	*None*	Maximum length of time that the ORB waits for second and subsequent message fragments before it times out. Set this property to zero to indicate no timeout. The default value is `30000`.

Steps for changing the location of the Object Request Broker

Perform the following steps to change the location of the Object Request Broker during installation:

1. Refer to Table 25 on page 153 for the default ORB-location information.
2. Change this default information during installation. In the Naming Server screen, the IBM WebSphere InterChange Server Installer prompts you for the IP address and port number for the IBM WebSphere InterChange Server instance. The installer saves this information in the appropriate variables of the product directory.

Table 25 also shows the variables within the CWSharedEnv file that specify the ORB location.

Table 25. Location of the IBM Java ORB

ORB location	IBM ORB property	Default value	CWSharedEnv variable
IP address	org.img.CORBA.ORBInitialHost	Name of the local host	ORB_HOST
Port number	org.omg.CORBA.ORBInitialPort	14500	ORB_PORT

Perform one of the following steps to change the location of the Object Request Broker *after* installation:

- Edit the appropriate variable (from the column in Table 25 labelled “CWSharedEnv variable”) in the CWSharedEnv file. Changing the ORB location within the CWSharedEnv file means that you change the ORB location for all invocations of all startup scripts that use the CWSharedEnv file. For example, to change the port number of the ORB to 15002, you can set the ORB_PORT property in the CWSharedEnv file, as follows:
  ```
  set ORB_PORT=15002
  ```
- Specify the appropriate IBM ORB property (from the column of Table 25 and its value in the command line of the component’s startup script with the -D command-line option. Changing the ORB location on the command line of a startup script means that you can change the ORB location only for that invocation of the component that the startup script is starting. For example, to change the port number of the ORB to 15002, you can specify the following -D option on the startup script’s command line:
  ```
  -Dorg.omg.CORBA.ORBInitialPort=15002
  ```

For information on how to change ORB properties, see “Steps for customizing the Object Request Broker” on page 151.

Using the IBM Transient Name Server

Using IBM Transient Naming Server (tnameserv) provides the naming service for the IBM WebSphere Business Integration system. When a component of the IBM WebSphere Business Integration system starts, it registers itself with the IBM Transient Naming Server. When the component needs access to another business-integration-system component, it uses the naming service to determine the information it needs to locate and start interacting with that component. For example, when an adapter needs to communicate with IBM WebSphere InterChange Server, it obtains the location of IBM WebSphere InterChange Server through the Transient Naming Server.

Note: The IBM Transient Naming Server is part of the IBM Java ORB. Therefore, it is installed automatically on the IBM WebSphere InterChange Server machine as part of the IBM WebSphere InterChange Server installation process.

Identification of registered components

The IBM WebSphere InterChange Server product provides the CosNameServer_Dump tool to list all valid IBM WebSphere InterChange Server ORB objects currently registered with the IBM Transient Naming Server. This tool is located in the bin subdirectory of the product directory. You invoke it with the following command: CosNameServer_Dump.bat.

Steps for using the Persistent Naming Server

When a component of the IBM WebSphere Business Integration system starts, it registers itself with the IBM Transient Naming Server and its CORBA object is stored in the memory of the Transient Naming Server. However, if the Transient Naming Server fails, its memory contents are lost. As a result, all components that had been registered with it must be rebooted so they can reregister with the naming service.

The Persistent Naming Server extends the capability of the IBM ORB Transient Naming Server so that the collection of CORBA objects that are registered with the Transient Naming Server are stored in a naming repository. The existence of the naming repository means that these CORBA references, rather than being only in the Transient Naming Server memory, are persistent; that is, they are available to other processes and IBM WebSphere InterChange Server components in the event that the Transient Naming Server fails. Other components do not need to shut down and restart in order to reregister with the naming service.

The default location of the naming repository is the following local file: `ProductDir\CxCosNameRepos.ior`.

Perform the following steps to change the location of the naming repository:

1. Edit the IBM WebSphere InterChange Server configuration file (InterchangeSystem.cfg).
2. Set the `CosNamingPersistencyFile` configuration parameter in the CORBA section. By default, the Persistent Naming Server is enabled; that is, references to CORBA objects are maintained in the naming repository.
3. For the naming server to run, you must explicitly start it with the `PersistentNameServer` startup file, located in the `bin` subdirectory of the product directory. This startup file takes the following steps:
   - Starts the IBM ORB Transient Naming Server
   - Starts the Persistent Naming Server to load the referenced CORBA objects into the naming repository.

   As part of its startup process, IBM WebSphere InterChange Server updates the naming repository by copying the CORBA objects currently registered with the Transient Naming Server into the naming-repository file. When each adapter starts, it updates the naming repository with its information. If IBM WebSphere InterChange Server has not yet started when the adapter starts, the naming repository is updated whenever IBM WebSphere InterChange Server does start.

   Note: If the Persistent Naming Server fails, you can restart it with the `PersistentNameServer` startup script. However, you do not need to restart IBM WebSphere InterChange Server or any started adapters.

Perform the following steps to turn off the Persistent Naming Server:

1. Edit the IBM WebSphere InterChange Server configuration file (InterchangeSystem.cfg).
2. Set the `CosNamingPersistency` configuration parameter (located in the CORBA section) to `false`.

Appendix A. WebSphere MQ reference

This appendix describes some of the commands that are used at the DOS command prompt to administer WebSphere MQ. A complete description of the WebSphere MQ commands can be found in the WebSphere MQ online documentation.

UNIX command	Windows command	What the command does
`start_mq`	`runmqsc` or `start_mq.bat`	Starts the WebSphere MQ command interpreter. The remaining commands can be entered after this command is run.
`end_mq`	`endmqm -i queue.manager`	Ends WebSphere MQ queues immediately. Must be run before deleting queues.
`clear_mq`	`clearorclear_mq.bat`	Clears WebSphere MQ queues.
`dltmqm queue.manager`	`dltmqm queue.manager`	Deletes WebSphere MQ queues.
`crtmqm queue.manager`	`crtmqm queue.manager`	Creates an WebSphere MQ queue manager.
`configure_mq` *<path to crossworlds_mq.tst>*	`configure_mq.bat` *<path to crossworlds_mq.tst>*	Configures WebSphere MQ queue manager.
`define`	`define`	Defines WebSphere MQ queues.

Appendix B. Requirements for restarting IBM WebSphere Business Integration system components

This appendix describes the restart requirements for administative tasks. The restart requirements for development or implementation tasks are described in the appropriate development guide or in the *System Implementation Guide*.

Use the following tables to decide if it is necessary to restart individual IBM WebSphere InterChange Server Components, including InterChange Server. A Dynamic listing in the Restart Requirements column means the component does not require a restart.

- "InterChange Server restart requirements"
- "Collaboration template restart requirements"
- "Collaboration object restart requirements" on page 158
- "Connector restart requirements" on page 158
- "Business object restart requirements" on page 159
- "Map restart requirements" on page 160
- "Relationship restart requirements" on page 160

InterChange Server restart requirements

Table 26 describes the restart requirements for InterChange Server.

Note: These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide.*

Table 26. Interchange Server restart requirements

Action	Restart requirement
Set Tracing options	Dynamic
Set Log/Trace file name	Dynamic
The max log/trace file size and number of archive files	Restart InterChange Server
Change InterChange Server name	Stop InterChange Server, recreate queues with new names, then start InterChange Server
Change InterChange Server password	Dynamic
Change repository database passwords	Dynamic
Add new class libraries (*.jar)	Add the name to the startup file, then restart InterChange Server
Change class libraries	Restart InterChange Server

Collaboration template restart requirements

Table 27 on page 158 describes the restart requirments for collaboration templates.

Note: These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide.*

Table 27. Collaboration template restart requirements

Action	Restart requirement
Recompile collaboration template without modifying the defined port names or adding or changing configuration property names.	Restart the collaboration object
Change port definition of the collaboration template.	Recompile the collaboration template; re-create the collaboration object
Change property names in the collaboration template.	Recompile the collaboration template; re-create the collaboration object

Collaboration object restart requirements

Table 28 describes the restart requirements for collaboration objects.

Note: These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide.*

Table 28. Collaboration object restart requirements

Action	Restart requirement
Add a new collaboration object	Dynamic, but the user must bind ports to connectors, then start the collaboration object
Delete a collaboration object	Requires that the collaboration object be stopped
Change port bindings	Collaboration object must be stopped before bindings can be changed then restarted after the change is made
Change configuration property	Dynamic
Change reuse of collaboration object	Restart the collaboration object
Change system trace level	Dynamic
Change collaboration trace level	Dynamic
Add e-mail notification recipient	Dynamic
Change number of concurrent events	Dynamic
Change user defined properties	Dynamic

Connector restart requirements

Table 29 describes the restart requirements for connectors.

Note: These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide.*

Table 29. Connector restart requirements

Action	Restart requirement
Add a new connector	Restart InterChange Server
Delete a connector	Dependent collaboration objects must be deleted first
Change configuration property (other than those mentioned below)	Restart connector component. See property description in Connector Configurator to determine action
Change supported business objects properties	Restart connector
Remove supported business object	Restart connector
Change trace level of the connector agent	Dynamic
Change trace level of the connector controller	Dynamic
Change poll frequency	Dynamic
Change controller store and forward mode	Dynamic

Table 29. Connector restart requirements (continued)

Action	Restart requirement
Change transport protocol	Restart InterChange Server
Change concurrent event triggered flows	Restart InterChange Server
Change poll start time	Restart connector
Change poll end time	Restart connector agent
Create new connector	Does not require a restart of InterChange Server
Recompile connector	Restart connector agent
New supported business objects	Restart connector agent
Bind maps to supported business object	Refresh connector screen in System Manager
Copy and paste connector	Restart InterChange Server. Until InterChange Server is restarted, the Associated Maps screen is blank because the information in this screen reflects the InterChange Server's run time only. Since this is a brand new connector with new supported BOs that are not yet reflected in InterChange Server's run time, the Associated Maps screen has no running BOs for binding maps.
Make changes to the values of the following standard properties: • `DeliveryTransport` • `AllowAnonymousConnections` • `EbGateway` • `RemoteWebGatewayURL` • `ListenPort` • `CACertificateDirectory` • `ConcurrentEvents` • `JMSBrokerName` • `JMSFactoryClassName` • `JMSUserNama` • `JMSPassword` • `NumberOfConnections`	From System Manager, right-click the server under Server Instances, then select Refresh. The Refresh option stops the connector, temporarily caches its transient state, removes the instance from memory and replaces it with the new instance, then starts the new instance.

Business object restart requirements

Table 30 describes the restart requirements for business objects.

Note: These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide.*

Table 30. Business object restart requirements

Action	Restart requirement
Add a business object	Dynamic
Delete a business object	Delete business object dependencies prior to deleting business object
Any change listed in this table.	If any tool, such as Map Designer or Process Designer, is connected to InterChange Server when any change listed in this table is made, that tool must be disconnected from, then reconnected, to the server.
Change the application-specific text for attributes.	Restart connector agent. Restarting InterChange Server is strongly recommended.
Change key attribute.	Restart connector agent.

Table 30. Business object restart requirements (continued)

Action	Restart requirement
Change default value.	If the connector uses the “UseDefaults” connector configuration property, the connector agent must be restarted. Restarting InterChange Server is strongly recommended.
Change max length for an attribute.	Restart connector agent.
Change required field for an attribute.	Restart connector agent.
Change trace level.	Dynamic
Change attribute names.	Restart connector agent.
Change business-object structure (removing or adding attributes/sub-objects).	Stop InterChange Server. Clean queues. Restart InterChange Server. The maps using the object must be updated and recompiled.

Map restart requirements

Table 31 describes the restart requirements for maps.

Note: These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide.*

Table 31. Map restart requirements

Action	Restart requirement
Add or update map definitions	User must compile the map, which reloads the map definition, then restart the map
Update map properties	Restart the map
Delete a map	The map must be stopped before it can be deleted
Recompile the map	Dynamic Compilation of maps puts them in active state by default, regardless of map instance reuse option.
Change reuse of map instances	Restart the map
Change trace level and data validation level	Restart the map
Add imported class libraries in the map	If the class library is not already included in the startup script of the server, it must be included and InterChange Server restarted.
Change imported class libraries	Restart InterChange Server If there are signature changes, the maps must also change and must be recompiled.

Relationship restart requirements

Table 32 on page 161 describes the restart requirements for relationships.

Note: These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide.*

Table 32. Relationships restart requirements

Action	Restart requirement
Create or change a relationship.	Stop the relationship before saving the changes with "Create runtime schema" option, and then restart the relationship.

Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Burlingame Laboratory Director
IBM Burlingame Laboratory
577 Airport Blvd., Suite 800

Burlingame, CA 94010
U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples may include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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

COPYRIGHT LICENSE

This information may contain sample application programs in source language, which illustrates 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.

Programming interface information

Programming interface information, if provided, is intended to help you create application software using this program.

General-use programming interfaces allow you to write application software that obtain the services of this program's tools.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

Warning: Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

Trademarks and service marks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries, or both:

IBM
the IBM logo
AIX
CICS
CrossWorlds
DB2
DB2 Universal Database
Domino
IMS
Informix
iSeries
Lotus
Lotus Notes
MQIntegrator
MQSeries
MVS
OS/400
Passport Advantage
SupportPac
WebSphere
z/OS

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

MMX, Pentium, and ProShare are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

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

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

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

The Monitor Definition Wizard and System Manager include software developed by the Eclipse Project (http://www.eclipse.org).

IBM WebSphere InterChange Server v4.3.0 and IBM WebSphere Business Integration Toolset, v4.3.0

Index

A

archive tables, backing up 107
automatic and remote restart of connectors 67

B

backing up
- archive tables 107
- collaboration class files 107
- components 106, 124
- repository 107

backup
- planning 105
- schedule 105

bar display option, default monitors 11
business object probe data file, deleting in System Monitor 21
business objects
- restart requirements 159

C

collaboration class files, backing up 107
collaboration objects
- configuring concurrent event-triggered flow processing 77
- configuring flow control 77
- configuring run-time properties 75
- management by SNMP agent 31
- operating 72
- restart requirements 158
- scheduling 101
- starting, stopping, and pausing 74
- states 72
- statistics 26

collaboration templates, restart requirements 157
community names
- SNMP Agent Configuration Manager 35

concurrent event-triggered flows, configuring for collaboration objects 77
configuring
- flow control 55, 69, 77
- persistent monitoring 22
- system-wide flow control 55

connecting to an InterChange Server instance 24
connectors
- automatic and remote restart 67
- commands for changing states 60
- configuring flow control 69
- enabling for MQ-triggered OAD 67
- initializing 59
- management by SNMP agent 31
- operating 57
- restart requirments 158
- restarting 66
- running, stopping, and pausing 59, 60
- scheduling 101
- shutting down 65
- starting manually 60
- states 57
- statistics 27

creating schedules 102
customizing the Object Request Broker (ORB) 151

D

database, changing password 51, 52
deleting schedules 104
disabling schedules 104
disconnecting from InterChange Server 88
displaying schedules 104

E

enabling schedules 104

F

Failed Event Manager 131
- checking access rights 140
- creating custom users and roles with Tomcat 134, 135
- logging on 136
- processing failed events 139
- viewing failed events 137

failed events, working with 131, 137, 139
flow control
- configuring for collaboration objects 77
- configuring for connectors 69
- configuring system-wide 55

Flow Manager
- finding unresolved flows 142
- handling unresolved flows 147
- managing queries 147
- starting 140
- viewing details of unresolved flows 145

H

High-Availability
- maintaining a Windows HA system 149
- managing 148
- supported environments 148

historical data, setting the frequency for capturing in System Monitor 19
historical statistics, deleting in System Monitor 21

I

IBM Transient Name Server 153
initializing a connector 59
installing
- Object Request Broker (ORB) 150
- SNMP Configuration Agent Manager 34

InterChange Server
- changing password 51
- configuring for flow monitoring 41
- connecting to an instance 24
- customizing startup parameters 49
- disconnecting from 88
- graceful shutdown 50

InterChange Server *(continued)*
- how it works with WebSphere Business Integration Monitor 41
- immediate shutdown 51
- management by SNMP agent 30
- operating 48
- password 51, 157
- restart requirements 157
- shutting down 50
- starting 48
- startup parameters 49
- statistics 25

InterChange Server Component Management view of System Manager 24

J

jobs, scheduling 99

L

line display option, default monitors 10
logging on to System Monitor 12
long-lived business processes 78

M

Management Information Base (MIB) 29
managing
- High-Availability (HA) systems 148
- Object Request Broker (ORB) 150
- run-time properties 147

maps
- operating 78
- restart requirements 160
- starting and stopping 80
- states 79

meter display option, default monitors 12
modifying schedules 103
Monitor Definition Wizard, creating new monitors 5
monitoring functions of InterChange Server Component Management view 24
monitors
- creating new 5
- customizing the visual appearance in System Monitor 18

monitors, default
- bar display option 11
- line display option 10
- meterdisplay option 12
- sample display options 9
- stacked bar display option 11
- table display option 9
- table tree display option 10
- using with System Monitor 2

N

new in release 4.0.0 x
new in release 4.0.1 x
new in release 4.1.0 x
new in release 4.1.1 x
new in release 4.2.0 viii
new in release 4.2.1 viii
new in release 4.2.2 vii

O

Object Request Broker (ORB)
- changing location 152
- customizing 151
- installing 150
- managing 150
- using the IBM Transient Name Server 153

opening relationships 89
operating
- collaboration objects 72
- connectors 57
- InterChange Server 48
- maps 78
- relationships 81

overriding schedules 102

P

password
- changing for database 51, 52
- changing for InterChange Server 51
- InterChange Server 51, 157

persistent monitoring 22
- accessing results 23
- configuring 22

persistent naming server 154

R

reconfiguring SNMP agent 32
refreshing System Manager 54
relationship data
- cleaning up participants 98
- filtering displayed participants 97
- printing relationship data 99
- searching for participants 96

Relationship Manager
- starting 86
- using 85

relationships
- copying participants 94
- creating relationship instances 92
- deactivating and activating participants 93
- deleting relationship instances 93
- loading and unloading business object files 94
- opening 89
- operating 81
- restart requirments 160
- retrieving relationship instances 89
- starting and stopping 82, 83
- states 81
- table caching 83
- working with 88
- working with relationship data 95

remote and automatic restart of connectors 67
repos_copy
- arguments added in release 4.2 113
- arguments removed in release 4.2 113
- command 108
- compiling components in the repository 116
- deleting components from the repository 117
- deploying a package to the repository 115
- exporting components to a package 118
- locale 119
- printing a list of components in the repository 119
- printing the command 114

repos_copy *(continued)*
- syntax 108
- usage scenarios 114
- validating the package 114
- validating the repository 116

Repository (InterChange Server)
- loading 108

repository, backing up 107
restart requirements 157
- business objects 159
- collaboration objects 158
- collaboration templates 157
- connectors 158
- InterChange Server 157
- maps 160
- relationships 160

restarting connectors 66
run-time properties, configuring for collaboration objects 75
running, stopping, and pausing connectors 59, 60

S

sample display options, default monitors 9
schedules
- creating 102
- deleting 104
- displaying 104
- enabling or disabling 104
- modifying 103
- overriding 102

scheduling
- collaboration objects 101
- connectors 101
- jobs 99

server access entries, configuring in SNMP Agent Configuration Manager 38
shutting down
- connectors 65
- InterChange Server 50
- InterChange Server gracefully 50
- InterChange Server immediately 51
- System Manager 54

SNMP agent
- connecting to from the SNMP Agent Configuration Manager 34
- how SNMP agent and SNMP manager communicate 30
- how SNMP works 28
  - community names 30
  - Management Information Base (MIB) 29
  - SNMP architecture 29
- how to use 31
- management of collaboration objects 31
- management of connectors 31
- management of InterChange Server 30
- reconfiguring 32
- starting and stopping 32
- using 28
- what SNMP manages 30

SNMP Agent Configuration Manager 33
- configuring community names 35
- configuring server access entries 38
- configuring trap forwarding 36
- connecting to the SNMP agent 34
- installing 34
- starting 34

SNMP manager, communiction with SNMP agent 30
stacked bar display option, default monitors 11
starting
- connectors manually 60
- Flow Manager 140
- InterChange Server 48
- Relationship Manager 86
- SNMP Agent Configuration Manager 34
- System Manager 53

starting and stopping
- maps 80
- relationships 82, 83
- SNMP agent 32

starting, stopping, and pausing
- collaboration objects 74

startup parameters
- customizing for InterChange Server 49
- InterChange Server 49

state change log, deleting in System Monitor 21
states
- capturing state changes in System Monitor 21
- collaboration objects 72
- connectors 57
- maps 79
- relationships 81
- reletionships 81

statistics
- collaboration objects 26
- connectors 27
- InterChange Server 25
- viewing and using from the InterChange Server Component Management view 24

System Manager
- InterChange Server Component Management view 24
- refreshing 54
- shutting down 54
- starting 53
- updating components 54
- using 53

System Monitor 13
- configuring existing views 16
- creating new monitors 5
- creating your own views 16
- deleting historical statistics 21
- deleting the business object probe data file 21
- deleting the state change log 21
- deleting views 17
- fine-tuning 17
  - capturing state changes 21
  - customizing the visual appearance of monitors 18
  - resetting run-time statistic values 20
  - setting the default view 18
  - setting the frequency for capturing historical data 19
  - setting the refresh rate for run-time values 19
- interface 13
- logging on 12
- using default monitors 2
- using display options 8
- using views 14

T

table caching, relationships 83
table display option, default monitors 9
table tree display option, default monitors 10
trap forwarding, configuring in SNMP Agent Configuration Manager 36
troubleshooting 131
typographic conventions vi

U

unresolved flows
  finding 142
  handling 147
  viewing details 145
updating components in System Manager 54

V

views
  configuring existing in System Monitor 16
  creating in System Monitor 16
  deleting in System Monitor 17
  setting the default in System Monitor 18
  using default in System Monitor 15
  using in System Monitor 14

W

WebSphere Business Integration Monitor 40
  administering 41
  configuring InterChange Server for flow monitoring 41
  how it works with InterChange Server 41
working with relationships 88

Printed in USA