

IBM WebSphere Business Integration Server Express
and Express Plus



System Administration Guide

Version 4.3

IBM WebSphere Business Integration Server Express
and Express Plus



System Administration Guide

Version 4.3

Note!

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

14May2004

This edition of this document applies to IBM WebSphere Business Integration Server Express, version 4.3, IBM WebSphere Business Integration Server Express Plus, version 4.3, 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.

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

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

© **Copyright International Business Machines Corporation 2004. All rights reserved.**

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

Contents

About this document	vii
Audience	vii
Related documents	vii
Typographic conventions	vii
New in this release.	ix
Chapter 1. Monitoring the system.	1
Overview of monitoring the system.	1
Using System Monitor	1
Setting up System Monitor	1
Steps for logging on to System Monitor	13
Overview of the System Monitor interface	15
Setting up views to monitor the system	16
Customizing data	18
Using persistent monitoring	23
Using System Manager to monitor the system.	25
Steps for connecting to an InterChange Server Express instance	25
Viewing and using statistics	26
Chapter 2. Administering components of the system	33
Overview of administering the system	33
Administering InterChange Server Express	34
Steps for starting InterChange Server Express	34
InterChange Server Express startup parameters	34
Shutting down InterChange Server Express.	35
Changing the InterChange Server Express and database passwords.	36
Using System Manager	38
Steps for starting System Manager.	39
Steps for shutting down System Manager	39
Steps for refreshing System Manager and updating components.	39
Steps for configuring system-wide flow control	40
Administering connectors	42
Viewing connector states	42
Starting, stopping, and pausing connectors	44
Steps for configuring flow control for connectors	51
Administering collaboration objects	52
Viewing collaboration object states.	52
Starting, stopping, and pausing collaboration objects	54
Configuring collaboration object run-time properties.	55
Administering maps	58
Viewing map states.	59
Starting and stopping maps	60
Administering relationships	60
Viewing relationship states	61
Starting and stopping relationships	62
Using Relationship Manager.	64
Working with relationships in Relationship Manager	68
Working with relationship data in Relationship Manager	74
Scheduling jobs in the system	78
Overview of scheduling jobs.	78
Steps for creating schedules	81
Steps for modifying schedules	82
Steps for deleting schedules	83
Steps for displaying schedules	83

Steps for enabling or disabling schedules	83
Backing up system components.	83
Backup schedule planning	84
Component backups	85
Using repos_copy	86
Repos_copy syntax	87
Repos_copy usage scenarios	91
Locale for repos_copy files	97

Chapter 3. Troubleshooting the system. 99

Using log and trace files for troubleshooting	99
Overview of log and trace files	99
Configuring e-mail notification of log messages	107
Configuring logging and tracing	110
Working with tracing in Log Viewer.	115
Steps for controlling server memory usage	121
Maintaining event archives	122
Managing Work-in Progress connections	123
Managing failures	123
Failure recovery for service calls	123
Strategies for InterChange Server Express recovery	124
Critical errors	126
Lost connection to application	127
Unknown connector agent status	127
Database connection failures	128
Flow failures	128
Working with failed events	129
Using Failed Event Manager	129
Using Flow Manager	136
Resolving collaboration deadlocks	143
Steps for detecting a collaboration deadlock	144
Steps for detecting group collaboration deadlocks	145
Steps for fixing a collaboration deadlock	145
Steps for preventing collaboration deadlocks	146
Managing IBM WebSphere MQ	147
Queue configuration errors	147
Failure to create new instances of MQQueueManager	147
Failure recovery and the queue manager	147
Memory checker segmentation fault	147
Steps for managing database connection pools	148
Steps for managing run-time properties	148
Managing the Object Request Broker	149
Object Request Broker installation	149
Steps for customizing the Object Request Broker	149
Steps for changing the location of the Object Request Broker	151
Using the IBM Transient Name Server	152

Appendix A. IBM WebSphere MQ reference 155

Appendix B. Requirements for restarting IBM WebSphere InterChange Server Express system components 157

InterChange Server Express 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

The products IBM^(R)WebSphere Business Integration Server Express and IBM^(R)WebSphere Business Integration Server Express Plus are made up of the following components: InterChange Server Express, the associated Toolset Express, CollaborationFoundation, and a set of software integration adapters. The tools in Toolset Express help you to create, modify, and manage business processes. You can choose from among the prepackaged adapters for your business processes that span applications. The standard processes template--CollaborationFoundation--allows you to quickly create customized processes.

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

Except where noted, all the information in this guide applies to both IBM WebSphere Business Integration Server Express and IBM WebSphere Business Integration Server Express Plus. The term WebSphere Business Integration Server Express and its variants refer to both products.

Audience

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

Related documents

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

You can download, install, and view the documentation at the following site:
<http://www.ibm.com/websphere/wbiserverexpress/infocenter>.

Note: Important information about this product may 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.

Typographic conventions

This document uses the following conventions:

<code>courier font</code>	Indicates information that you type, such as a command name, file name, information that you type, or information that the system prints on the screen.
<i>italic, italic</i>	Indicates a new term the first time it appears, words defined in the text, or a cross-reference.
bold	Indicates a GUI element, such as a radio button, field name, or folder name.

<i>blue outline</i>	A blue outline, which is visible only when you view the manual online, indicates a cross-reference hyperlink. Click inside the outline to jump to the object of the reference.
{ }	In a syntax line, curly braces surround a set of options from which you must choose one and only one.
[]	In a syntax line, square brackets surround an optional parameter.
...	In a syntax line, ellipses indicate a repetition of the previous parameter. For example, <code>option[...]</code> means that you can enter multiple, comma-separated options.
< >	In a naming convention, angle brackets surround individual elements of a name to distinguish them from each other, as in <code><server_name><connector_name>tmp.log</code> .
/, \	In this document, backslashes (\) are used as the convention for directory paths. For UNIX installations, substitute slashes (/) for backslashes. All IBM WebSphere InterChange Server Express product path names are relative to the directory where the IBM WebSphere InterChange Server Express product is installed on your system.
<i>%text%</i>	Text within percent (%) signs indicates the value of the Windows text system variable or user variable.
<i>ProductDir</i>	Represents the directory where the product is installed.

New in this release

This is the first release of this guide.

This is the first release of this guide.

Chapter 1. Monitoring the system

Monitoring the overall health of the IBM WebSphere InterChange Server Express system includes monitoring all InterChange Server Express components, such as connectors and collaboration objects, and the connection to all integrated applications. You can monitor the system using capabilities in System Monitor and System Manager.

Overview of monitoring the system

This chapter covers all methods for system monitoring using System Monitor and System Manager tools and includes the following topics:

“Using System Monitor”

“Using System Manager to monitor the system” on page 25

Using System Monitor

System Monitor is a tool that allows you to monitor the IBM WebSphere InterChange Server Express system from the Web. It allows you to configure how you view the data and also allows you to view historical data and current data. System Monitor uses *monitors*, definitions of the information you want to access when monitoring the system, and *views*, definitions of the monitors that you want displayed in your Web browser. Monitors and views allow you to customize how you monitor the system. You determine the monitors you use before you begin using System Monitor (see “Setting up System Monitor”). You determine the views you want to use when you are monitoring the system (see “Setting up views to monitor the system” on page 16).

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 33.

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 13

“Overview of the System Monitor interface” on page 15

“Setting up views to monitor the system” on page 16

“Customizing data” on page 18

“Using persistent monitoring” on page 23

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 Express versions 5.0.2 or 5.1, IBM WebSphere Application Server versions 5.0.2 or 5.1, or Tomcat versions 4.1.24 or 4.1.27 (using IBM JDK 1.3.1)	DB2 Client (if using DB2 for repository database)	<ul style="list-style-type: none"> 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

Perform the following step to gain a better understanding of what information is included in each monitor:

- Refer to Table 2 below to determine if you want to use the default monitors provided with the System Monitor. (Table 3 on page 9 contains a description of display options listed in Table 2. and “Examples of display options” on page 10 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)	<ul style="list-style-type: none"> 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

Table 2. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor
Collaboration Statistics	Current status and statistics of all collaborations in the system: <ul style="list-style-type: none"> • Status • Start time • Total flows • Successful flows • Failed flows • Events in process • Queued events • Max concurrent events 	Table	<ul style="list-style-type: none"> • Start, stop, pause, and shut down
Connector Statistics	Current status and statistics of all connectors: <ul style="list-style-type: none"> • Status • Start time • Total up time • Business objects received • Business objects sent • Agent status 	Table	<ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • 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

Table 2. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor
Connector Subscriptions	Current statistics of the subscriptions for a particular connector: <ul style="list-style-type: none"> • Collaboration object • Initiator 	Table	None
Collaboration Events	Current statistics of collaboration events, which includes the following information: <ul style="list-style-type: none"> • Events in process • Queued events 	Bar	None
Historical Server Statistics	Server statistics for a specific period of time. Statistical information: <ul style="list-style-type: none"> • Successful calls • Failed calls • Total calls • Successful events • Failed events • Total events • Successful flows • Failed flows • Total flows Time intervals: <ul style="list-style-type: none"> • 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: <ul style="list-style-type: none"> • Successful flows • Failed flows • Total flows Time intervals: <ul style="list-style-type: none"> • 15 min., 30 min., hourly, daily, weekly, or monthly • Start date • End date 	<ul style="list-style-type: none"> • Line • Stacked bar • Bar 	None

Table 2. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor
Historical Collaboration Flows Stack	<p>Flow statistics of a particular collaboration for a specific period of time at certain time intervals.</p> <p>Statistics information:</p> <ul style="list-style-type: none"> • Successful flows • Failed flows • Total flows <p>Time intervals:</p> <ul style="list-style-type: none"> • 15 min., 30 min., hourly, 4 hours, 12 hours, daily, weekly, or monthly • Start date • End date 	Stacked bar	None
Historical Collaboration Flows Line	<p>Flow statistics of a particular collaboration for a specific period of time at certain time intervals.</p> <p>Statistics information:</p> <ul style="list-style-type: none"> • Successful flows • Failed flows • Total flows <p>Time intervals:</p> <ul style="list-style-type: none"> • 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	<p>Current state of collaboration objects and connectors under Flow Control:</p> <ul style="list-style-type: none"> • 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

Table 2. Default monitors (continued)

Default monitor	Definition	Display options	Available operations when viewing monitor
State Change Log	<p>Current persisted state changes on a component for a specified time period.</p> <p>State change information:</p> <ul style="list-style-type: none"> • Time stamp • State <p>Time intervals:</p> <ul style="list-style-type: none"> • Start date • End date 	Table	None

Steps for creating additional monitors

Before you begin: 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 Express instance is running.

4. Right-click the server instance, then click **Monitor Definition Wizard**. The Monitor Definition Wizard appears. See Figure 1 on page 7..




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 9, and for examples of how the display options appear in System Monitor, see “Examples of display options” on page 10.

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 on page 8).




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**.




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 6) or using monitors in System Monitor (see “Steps for setting display properties” on page 19):

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 10.)

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
<ul style="list-style-type: none"> • Table • Table tree 	<ul style="list-style-type: none"> • Columns to display • Order of columns • Number of rows to display 	<ul style="list-style-type: none"> • Font and color settings of the labels and data • Number of rows to display
<ul style="list-style-type: none"> • Stacked bar • Line • Bar 	None	<ul style="list-style-type: none"> • Font and color settings of the labels and data • Show or hide values

Table 3. Configurable display options for monitors (continued)


Display option	Properties you can configure when building monitors in Monitor Definition Wizard	Properties you can configure when using monitors in System Monitor
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 Express system.

Connector Statistics 

Display items

1-9 of 9 items

Select all displayed items

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

Sat Mar 15 16:16:21 PST 2003

Figure 4. Table display option

System Overview



Display items

<<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

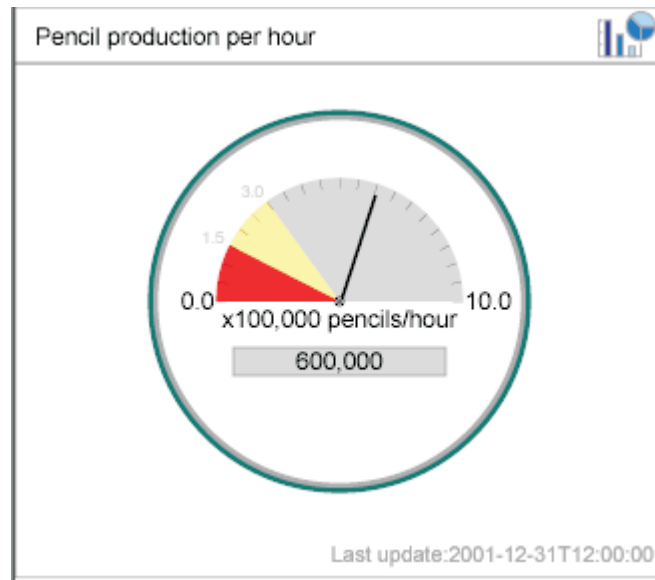


Figure 9. Meter display option

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 Express on the machine being monitored.
- Make sure System Monitor and the application server are installed. For installation instructions, see *WebSphere Business Integration Server Express 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 Express.

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, WebSphere Application Server Express, or Tomcat. Refer to the *WebSphere Business Integration Server Express Installation Guide* for additional information about setting up System Monitor to work with WebSphere Application Server, WebSphere Application Server Express, or Tomcat.
 - If System Monitor is installed as part of WebSphere Application Server Express installation, the URL is:
`http://HostName:7089/ICSMonitor`

 where *HostName* is the host name of the Web server machine.
 - If you used the installation script with a “y” in the 5th parameter, the URL is:
`http://HostName/ICSMonitor`

where *HostName* is the host name of the Web server machine.

- If you used the installation script with a “n” in the 5th parameter and a port number in the 6th parameter, the URL is:

`http://HostName:nnnn/ICSMonitor`

where *HostName* is the host name of the Web server machine and *nnnn* is the port number.

- 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 Express System Monitor login window appears (see Figure 10).

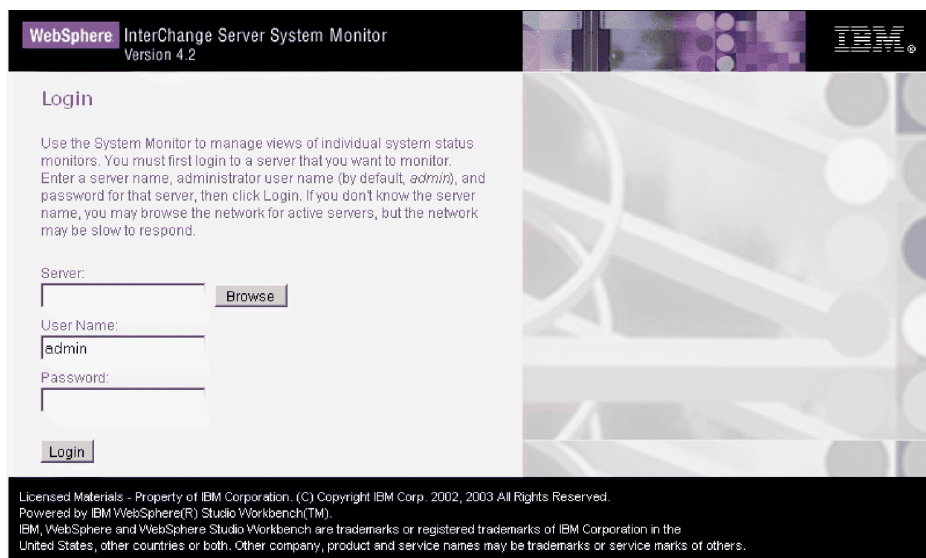


Figure 10. System Monitor, Login window

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




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

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 17), which allows you to create, configure, or delete views. It also allows you to set the default view you see when you log on to System Monitor. For instructions on creating, configuring, or deleting views, or setting the default view you see when you log on, refer to the following sections:
 - “Steps for creating your own views” on page 17
 - “Steps for configuring views” on page 18
 - “Steps for deleting views” on page 18
 - “Steps for setting a default view” on page 19
- **Set Options link:** The **Set Options** link opens the Set Options dialog box (see Figure 14 on page 20), which allows you to 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 20
- “Steps for setting the frequency for historical data capture” on page 20
- “Steps for resetting run-time statistic values” on page 21

- “Steps for capturing state changes” on page 22
- “Steps for deleting the state change log” on page 22
- “Steps for deleting historical statistics” on page 22
- “Steps for deleting the business object probe data file” on page 22
- **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 19.
- **Logoff link:** The **Logoff** link allows you to log off System Manager.
- **Help link:** The **Help** link opens an HTML page with the following information:
 - A link to download the documentation set for the IBM WebSphere InterChange Server Express 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 contain 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 13.

This section covers the following topics:

“Steps for using default views”

“Steps for creating your own views” on page 17

“Steps for configuring views” on page 18

“Steps for deleting views” on page 18

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	<ul style="list-style-type: none"> • Collaboration Events monitor displayed in bar chart, and • Event Rate monitor displayed in a meter

Table 4. Default views (continued)

Default view	Monitor(s) and display options
Collaboration History	<ul style="list-style-type: none"> 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	<ul style="list-style-type: none"> Business Objects monitor displayed in a table Connector Subscriptions monitor displayed in a table
Maps and Relationships	<ul style="list-style-type: none"> Map Status monitor displayed in a table Relationship Status monitor displayed in a table
Server Statistics	<ul style="list-style-type: none"> Server Statistics displayed in a stacked bar chart Database Connections displayed in a table Message Queues displayed in a table
Server History	<ul style="list-style-type: none"> 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 10.

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).

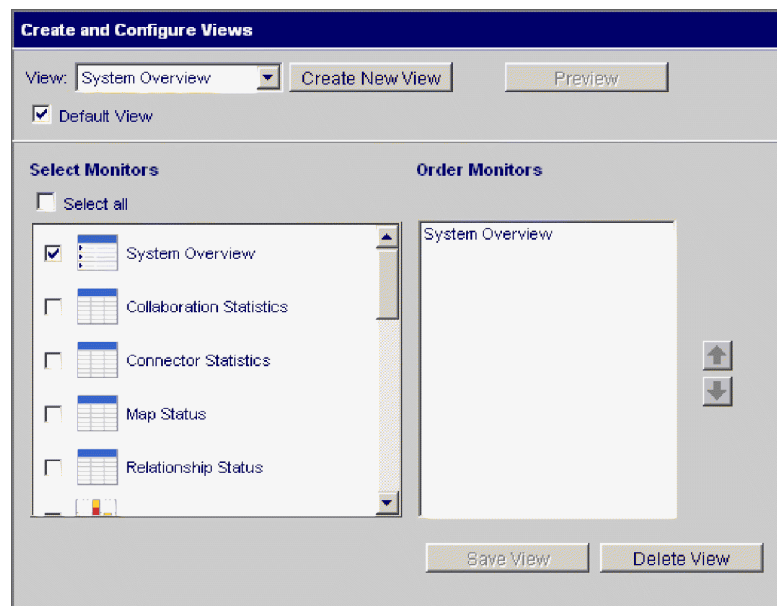


Figure 12. System Monitor, Create and Configure Views dialog box

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

3. Type a name for the view in the **View Name** field, then click **OK**. The new view name appears in the **View** field of the Create and Configure Views dialog box.
4. Select one or more monitors 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 on page 17).
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 17).
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.

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 19

“Steps for setting display properties” on page 19

“Steps for setting the refresh rate for run-time values” on page 20

“Steps for setting the frequency for historical data capture” on page 20

“Steps for resetting run-time statistic values” on page 21

“Steps for capturing state changes” on page 22

“Steps for deleting the state change log” on page 22

“Steps for deleting historical statistics” on page 22

“Steps for deleting the business object probe data file” on page 22

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 17).
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

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

1. 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.

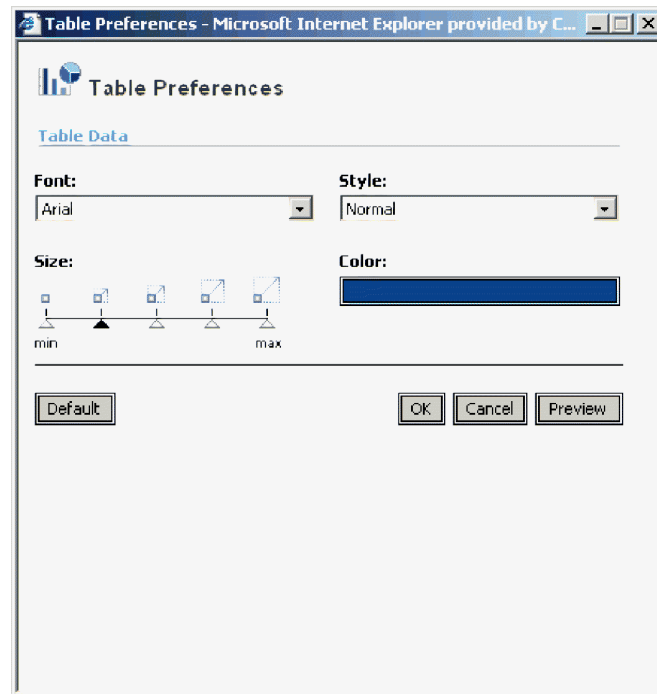


Figure 13. System Monitor, Table Preferences dialog box

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 9 and 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).

The screenshot shows the 'Set Options' dialog box with the following sections and controls:

- How often do you want to update statistics?**
 - Refresh Rate: 0 seconds
 - Submit button
- How frequently should historical data be captured?**
 - Component Type: Server
 - Frequency: NONE
 - Review all interval settings link
 - Submit button
- Do you want to reset component statistics?**
 - Component Type: Server (all components)
 - Submit button
- Do you want to capture state changes of a particular component?**
 - Component Type: Server
 - Capture State Changes checkbox
 - Submit button
- Do you want to delete the state change log for all components?**
 - Start Date: [calendar icon]
 - End Date: [calendar icon]
 - Delete button
- Do you want to delete the historical statistics for all components?**
 - Start Date: [calendar icon]
 - End Date: [calendar icon]
 - Delete button
- Do you want to delete the data for a Business Object Probe?**
 - Business Object: JTextPT:BOProbe.LinkTransition_46.LinkTransition_46
 - Start Date: [calendar icon]
 - End Date: [calendar icon]
 - Delete button

Figure 14. System Monitor, Set Options dialog box

2. 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 21).

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
Submit Changes	

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 **Submit**.

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 20).
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 20).
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 20).
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 20).
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 20).
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 Express 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 18:

- “Steps for setting the frequency for historical data capture” on page 20
- “Steps for resetting run-time statistic values” on page 21
- “Steps for capturing state changes” on page 22
- “Steps for deleting the state change log” on page 22
- “Steps for deleting historical statistics” on page 22
- “Steps for deleting the business object probe data file” on page 22

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 24).




Figure 16. Edit Configuration tool, Misc tab

3. Under **Persistent Monitoring**, do the following:
 - If you want InterChange Server Express 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 Express 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 16. 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 17.

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 Express 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 Express 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 33.

The following section describes how to work in System Manager to connect to an InterChange Server Express 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 Express instance.

Steps for connecting to an InterChange Server Express instance

Perform the following steps to connect to an InterChange Server Express instance.

1. Open System Manager.
2. In the InterChange Server Component Management view, the lower left quadrant of System Manager, right-click the InterChange Server Express instance to which you want to connect, and click **Connect** on the pop-up menu (see Figure 17).




Figure 17. Connect to InterChange Server Express

Note: If you want to register a new InterChange Server instance, right-click and click **Register Server**.

3. 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 Express instance** icon changes from red to green, and any objects that have been deployed to that server appear in folders beneath the server (see Figure 18 on page 26).

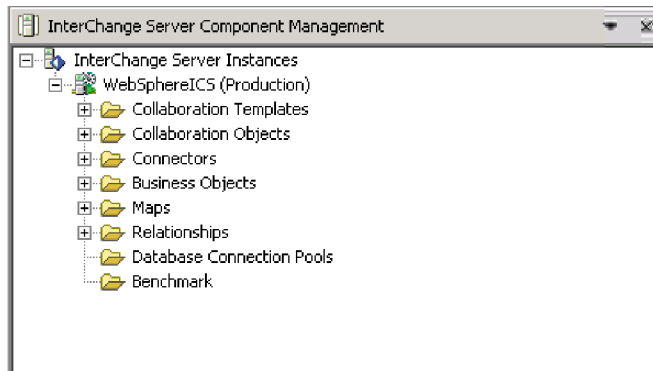


Figure 18. InterChange Server Component Management view, showing connection to InterChange Server Express

Viewing and using statistics

You can monitor statistics for the InterChange Server Express environment in the InterChange Server Component Management view to help you better manage the system. Statistics can be viewed for InterChange Server Express, 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 procedures describe how to use the Server Statistics window, the Collaboration Object Statistics window, and the Connector Statistics window:

- "Steps for viewing server statistics"
- "Steps for viewing collaboration object statistics" on page 28
- "Steps for viewing connector statistics" on page 30

Note: Before you can see any system statistics, System Manager must be connected to an InterChange Server Express instance. For instructions on connecting to an InterChange Server Express instance, see "Steps for connecting to an InterChange Server Express instance" on page 25.

Steps for viewing server statistics

Perform the following steps to check InterChange Server Express statistics:

1. Open System Manager and go to the InterChange Server Component Management view.
2. Right-click the InterChange Server Express instance and click **Statistics**.
The server statistics appear in the upper-right quadrant of System Manager (see Figure 19 on page 27).

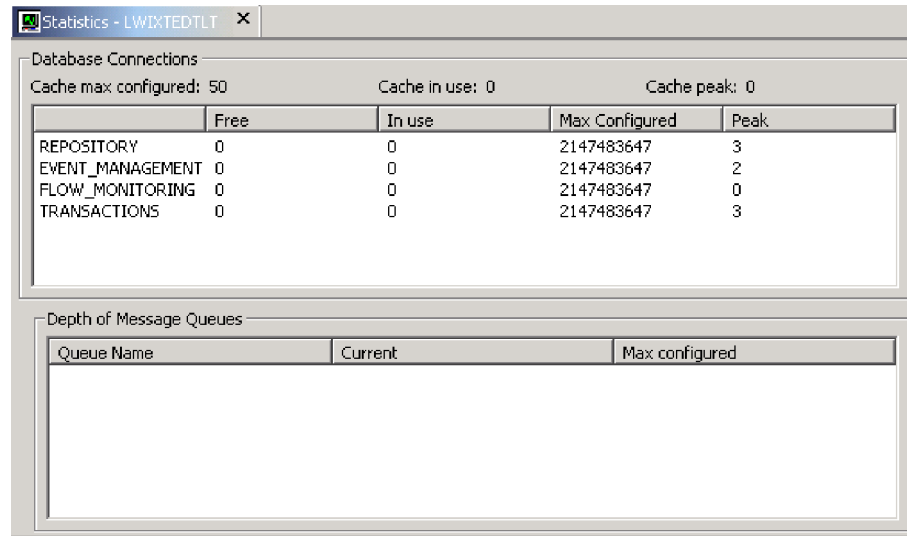


Figure 19. Statistics window for InterChange Server Express

3. Under **Database Connections** you can view the following information:

- Find out how many database connections the InterChange Server Express system's connection cache is currently using and the peak amount used since the server was booted. This can help you tune the interaction with the underlying Database server for InterChange Server Express. 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.
- Look for the connection pool that is consuming the most number of connections. This can help you configure InterChange Server Express 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 Express and the database management system (DBMS). For information about these parameters, see *WebSphere Business Integration Server Express Installation Guide*.

The Database connections area shows statistics for:

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

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

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

In use	The current number of connections used by this connection pool.
Max Configured	The maximum number of connections configured. This is the value in the InterchangeSystem.cfg file for the attribute MAX_CONNECTIONS in the respective subsections of the different connection pools (Event Management, Transactions, Repository). If this attribute is not configured, it displays Default.
Peak	The highest number of connections used by the server from this pool since the server was booted.
4. Under Depth of Message Queues ,	you can view a list of all the subscription queues in the configured queue manager. The following statistics are displayed:
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.

Steps for viewing collaboration object statistics

Perform the following steps to check collaboration object statistics:

1. Open System Manager and go to the InterChange Server Component Management view.
2. Expand the Collaboration Objects folder.
3. Right-click the collaboration object for which you want to view statistics and click **Statistics**.

The statistics for that collaboration object appear in the upper-right quadrant of System Manager (see Figure 20 on page 29).

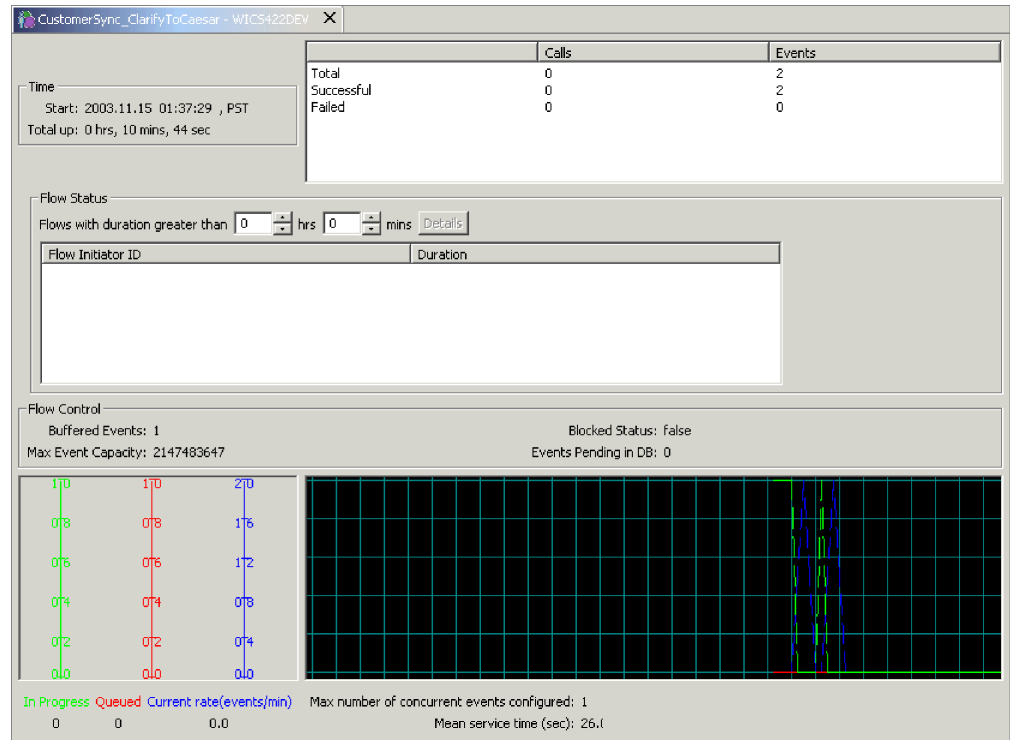


Figure 20. Statistics window for collaboration objects

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

4. 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, because 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.
5. To view general statics, do the following in the top section of the statistics window:
 - 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*
6. Under **Flow Status**, you can do the following:
 - 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.
7. Under **Flow Control**, you can do the following:

- 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 57
8. At the bottom of the statistics window, you can do the following:
- 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.

Steps for viewing connector statistics

Perform the following steps to check connector statistics:

1. Open System Manager and go to the InterChange Server Component Management view.
2. Expand the Connectors folder.
3. Right-click the connector for which you want to view statistics and click **Statistics**.

The statistics for that connector appear in the upper-right quadrant of System Manager (see Figure 21 on page 31). The window for provides information about the running connector: 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 flow control information.




Figure 21. Statistics for connectors

4. Under **Time**, you can view the start time of the connector and how long it has been running.
5. Under **Business Objects**, you can view the total number of business objects received and sent during the time the connector has been running.
6. Under **Business Objects Sent and Received**, you can view the names of the business objects that 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.
7. Under **Subscriptions**, you can view a list of 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.
8. Under **Flow Control**, you can view the number of buffered events and the number of events pending in the database. Two configurable Flow Control properties are listed: **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 51.

Chapter 2. Administering components of the system

This chapter describes some of the tasks you may need to perform when administering an IBM WebSphere InterChange Server Express system. For instructions on starting InterChange Server Express for the first time, refer to *WebSphere Business Integration Server Express Installation Guide*.

This chapter covers the following topics:

“Overview of administering the system” on page 33

“Administering InterChange Server Express” on page 34

“Using System Manager” on page 38

“Administering connectors” on page 42

“Administering collaboration objects” on page 52

“Administering maps” on page 58

“Administering relationships” on page 60

“Scheduling jobs in the system” on page 78

“Backing up system components” on page 83

“Using repos_copy” on page 86

Overview of administering the system

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

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

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 WebSphere InterChange Server Express is started. The only required order for starting components is if you are using IBM Java Object Request Broker (ORB) located on a different machine than WebSphere InterChange Server Express, you must start the ORB before starting WebSphere InterChange Server Express.

When shutting down an instance of WebSphere InterChange Server Express, 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 InterChange Server Express system, make sure all of the necessary third-party software is running. This includes the database on which the InterChange Server Express 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 *WebSphere Business Integration Server Express Installation Guide*.

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

The following tasks describe 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 Express is located, start the ORB. For more information about using ORB, see “Managing the Object Request Broker” on page 149.
3. If you are using IBM WebSphere MQ, start the MQ Listener.
4. Start InterChange Server Express. See “Steps for starting InterChange Server Express” on page 34.
5. Start any connectors not automatically started. See “Starting, stopping, and pausing connectors” on page 44.
6. Start collaborations. See “Viewing collaboration object states” on page 52.
7. Start System Manager. See “Using System Manager” on page 38.

Administering InterChange Server Express

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

“Steps for starting InterChange Server Express” on page 34

“InterChange Server Express startup parameters”

“Shutting down InterChange Server Express” on page 35

“Changing the InterChange Server Express and database passwords” on page 36

Steps for starting InterChange Server Express

Perform the following step to start InterChange Server Express:

- Click **Start > Programs > IBM WebSphere Business Integration Express > InterChange Server Express > InterChange Server Express**.

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

InterChange Server Express startup parameters

To customize the InterChange Server Express startup parameters, modify the InterChange Server Express shortcut or the `start_server.bat` file.

The parameters in table Table 5 customize the startup of InterChange Server Express.

Table 5. InterChange Server Express startup parameters

Parameter	Function
-c <i>configFile</i>	Name of the configuration file to be used during startup. The default is InterchangeSystem.cfg.
-design	Specifies that the server is started in the design mode.
-i	Allows InterChange Server Express to start up and ignore all error messages.
-p <i>password</i>	Specifies the password to access InterChange Server Express. 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 <i>serverName</i>	Specifies the name of the InterChange Server Express instance. The name is case-sensitive.
-u <i>loginName</i>	Specifies the user login name for InterChange Server Express. 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	Displays the version of InterChange Server Express, then exits.

Shutting down InterChange Server Express

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

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

Steps for shutting down InterChange Server Express gracefully

Perform the following steps to shut down the server gracefully:

1. In System Manager, go to the InterChange Server Component Management view.
2. Right-click the InterChange Server Express instance and click **Shut Down > Gracefully**.

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

If you want 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 Express shuts down.

Steps for shutting down InterChange Server Express immediately

Perform the following steps to shut down the server immediately:

1. In System Manager, go to the InterChange Server Component Management view.
2. Right-click the InterChange Server Express instance and click **Shut Down > Immediately**.

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 “Flow failures” on page 128 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 InterChange Server Express 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 Express and database passwords

Password encryption provides a measure of security for protecting the IBM WebSphere InterChange Server Express system and underlying databases from unauthorized user entry. The encrypted string for each of the passwords is stored in InterChange Server Express 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 InterChange Server Express 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 Express password or the database password in System Manager.

The InterChange Server Express user name and password are required during repository copy and restoration. See “Using repos_copy” on page 86.

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

“Steps for changing the InterChange Server Express password”

“Steps for changing the database passwords” on page 37

Steps for changing the InterChange Server Express password

Perform the following steps to change the password for InterChange Server Express:

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

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

Attention: The InterChange Server Express 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 Express does not start.

Steps for changing the database passwords

The repository database passwords can be changed through System Manager after the InterChange Server Express system is running.

Perform the following steps to change the database passwords:

1. In the InterChange Server Component Management view, right-click the InterChange Server Express instance whose database password you want to change, then click **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 22 on page 38).

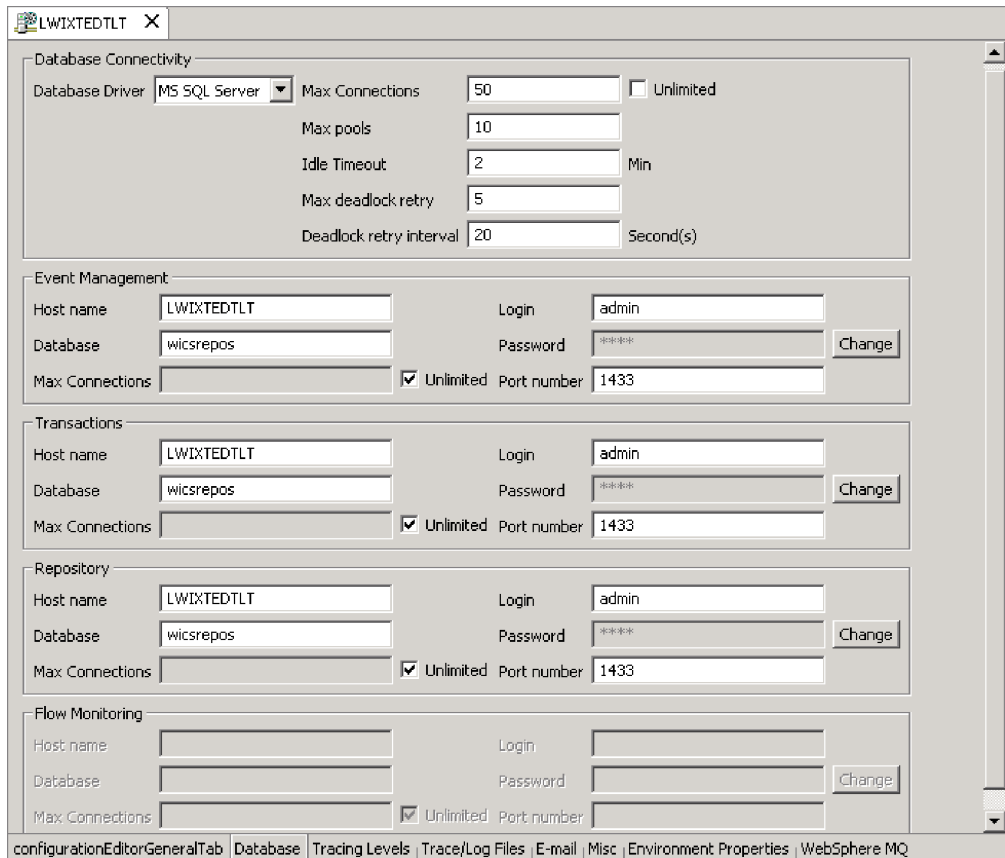


Figure 22. Database tab of Edit Configuration window

3. Change any of the database passwords by doing the following:
 - In the section for the appropriate database (**Event Management**, **Transactions**, or **Repository**), click **Change**.
A dialog box for changing the password appears.
 - Type the current password in the **Old Password** field.
 - Type a new password in the **New Password** field.
A maximum of 30 characters is allowed.
 - Type the new password again 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” on page 39

“Steps for shutting down System Manager” on page 39

“Steps for refreshing System Manager and updating components” on page 39

“Steps for configuring system-wide flow control” on page 40

Steps for starting System Manager

Perform the following step to start System Manager:

- Click **Start > Programs > IBM WebSphere Business Integration Express > Toolset Express > Administrative > System Manager**.

The System Manager perspective of the IBM WebSphere Studio Workbench appears (see Figure 23).

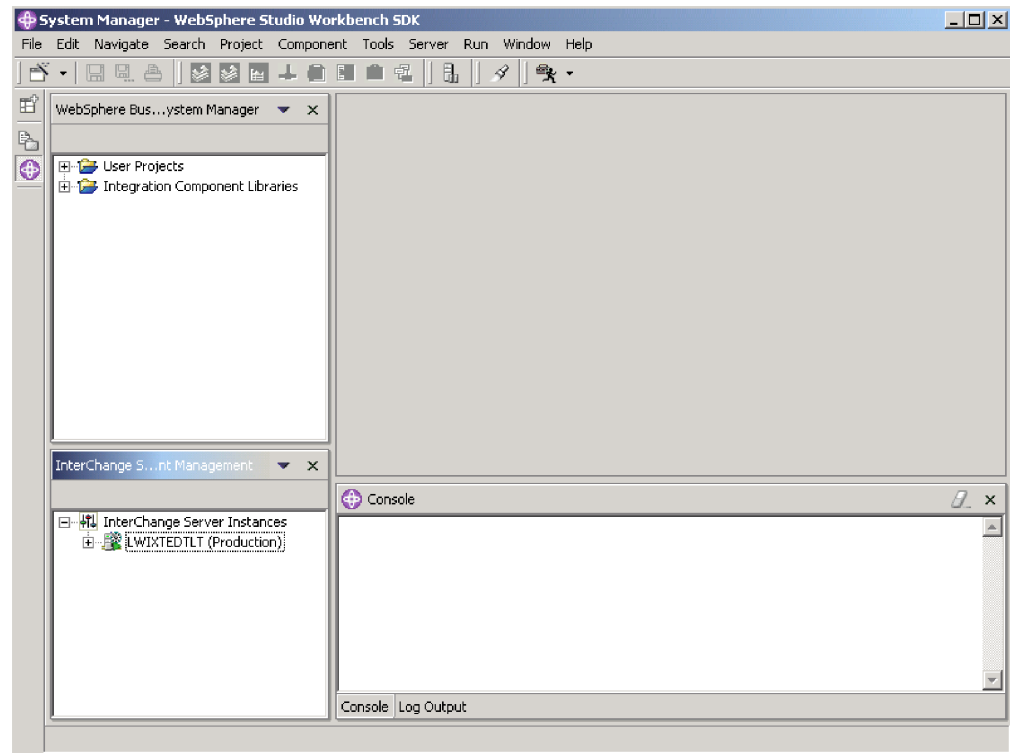


Figure 23. System Manager

Steps for shutting down System Manager

Perform the following step to shut down System Manager:

- In IBM WebSphere Studio Workbench, click **File > Exit**.

Note: Be sure to shut down any InterChange Server Express instances before shutting down System Manager. For instructions on shutting down InterChange Server Express, see “Shutting down InterChange Server Express” on page 35.

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 Express. 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 Express

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 steps to refresh System Manager:

1. In System Manager, right-click the server instance.
2. Click **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 Express 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 51 or "Steps for configuring flow control for collaboration objects" on page 57.

Note: Configuration changes for individual connectors or collaboration objects are dynamic, meaning they do not require InterChange Server Express to be rebooted. System-wide configuration changes for flow control require InterChange Server Express 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 System Manager in the InterChange Server Component Management view. 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 16. For more information on viewing the flow control from the InterChange Server Component Management view, see "Steps for viewing collaboration object statistics" on page 28 or "Steps for viewing connector statistics" on page 30.

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

1. In the InterChange Server Component Management view, right-click the InterChange Server Express instance for which you want to configure flow control, and click **Edit Configuration**. The upper-right quadrant of System Manager changes to an editing tool in which many system properties can be changed.

- Click the **Misc** tab. A dialog box appears with a **Flow Control** section (see Figure 24).

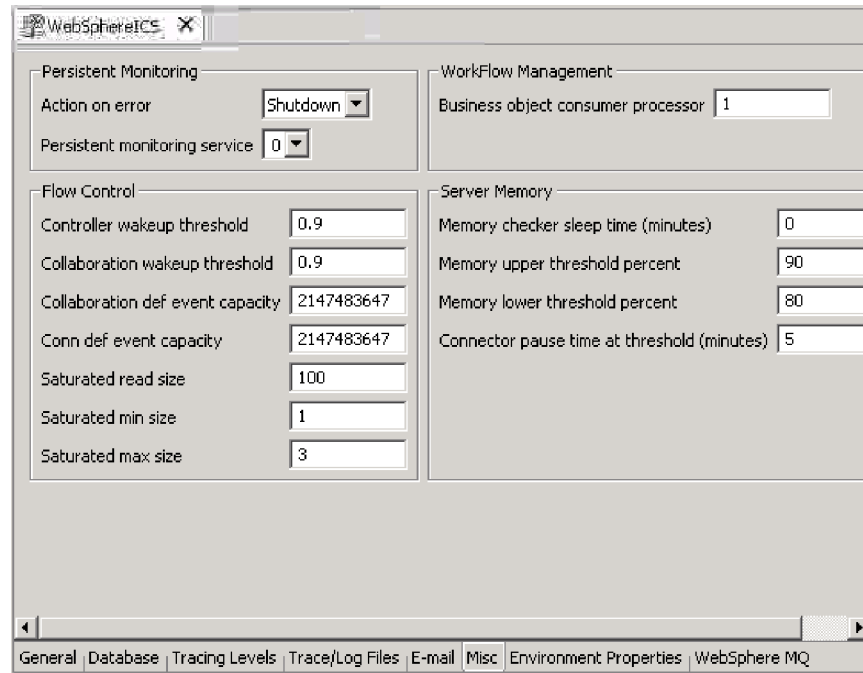


Figure 24. Edit configuration tool, Misc tab

- Under **FlowControl**, type information in the following fields:

-
-
- **Controller wakeup threshold:** This property applies to connector event queues. It has a decimal value ranging from 0 to 1, but not including 0 or 1. Connector event queues are always of the blocking type, meaning that if the queue is full, they do not allow new events to be added. After a queue becomes full, the connector becomes blocked. When the queue size equals or falls below the value of the connector wakeup threshold multiplied by the maximum event capacity of that connector ($\text{CONTROLLER_WAKEUP_THRESHOLD} \times \text{MaxEventCapacity}$), the connector becomes reactivated.
- **Collaboration wakeup threshold:** This property applies to collaboration object event queues. It has a decimal value ranging from 0 to 1, but not including 0 or 1. This property applies only to blocking-type collaboration objects, meaning that it does not allow the connector to add more events to the collaboration queue. When the queue size equals or falls below the value of the collaboration object wakeup threshold multiplied by the maximum event capacity of that connector ($\text{COLLABORATION_WAKEUP_THRESHOLD} \times \text{MaxEventCapacity}$), the connector is able to add more events to the collaboration queue for processing.
- **Collaboration def event capacity:** 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.
- **Conn def event capacity:** 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.

- **Saturated read size:** 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.
 - **Saturated min size:** 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.
 - **Saturated max size:** This property applies to saturated readers, which are readers that process saturated events in the database, then add those events to the appropriate collaboration object queue. This property reflects the maximum number of threads doing these activities. The default is 3.
4. Save the changes you made to the InterChange Server Express configuration by clicking **File > Save <ServerName>**.
 5. Restart InterChange Server Express.

Administering connectors

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

You can start, pause, stop, and shut down connectors from either System Monitor or System Manager.

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 in System Manager. To use System Manager, follow the instructions in “Steps for connecting to an InterChange Server Express instance” on page 25.

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

Steps for using System Monitor to view connector states

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

1. Log on to System Monitor (see “Steps for logging on to System Monitor” on page 13).
2. If the System Overview view is not displayed, click the **System Overview** link under **Views**. The System Overview Monitor appears (see Figure 11 on page 15).

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 16 for instructions.

3. Click the triangle next to the name of the server to display a list of components on the system.
4. Click the triangle next to a running collaboration to display its associated connectors (see Figure 25 on page 43).




Figure 25. System Monitor, System Overview displaying connector state

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 Express instance that contains the connector you want to view. See “Steps for connecting to an InterChange Server Express instance” on page 25 for instructions on connecting to an InterChange Server Express instance.
2. Expand the InterChange Server Express instance, then expand the Connectors folder.

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

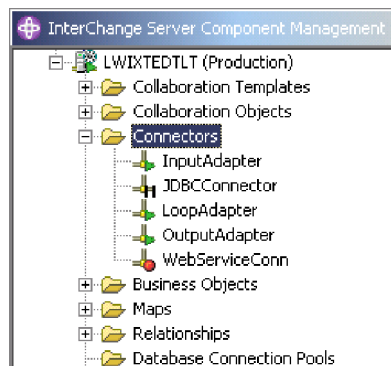


Figure 26. Connectors folder in the InterChange Server Component Management view

Table 6 lists the connector states represented by the display color of each connector and shows what actions are being performed during that state.

Table 6. Connector states

Connector state (color)	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.

This section covers the following topics:

“Connector initialization” on page 44

“Steps for starting, stopping, and pausing connectors” on page 44

“System Manager commands for changing connector states” on page 45

“Steps for manually starting a connector” on page 45

“Shutting down a connector” on page 47

“Restarting a connector” on page 48

“Steps for setting automatic and remote restart for a connector” on page 48

“Steps for using OAD as a Windows service” on page 50

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 manually starting a connector” on page 45.

If the connector does not start, check to make sure that the command line to start it includes the current InterChange Server Express name. For more information, refer to *WebSphere Business Integration Server Express Installation Guide*.

Steps for starting, stopping, and pausing connectors

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

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

1. When viewing the System Overview view (see Figure 25 on page 43), specify a connector by selecting the check 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 27 on page 45).



Figure 27. System Monitor, icons for starting, pausing, restarting, or shutting down components

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

1. In System Manager, go to the Connectors folder in the InterChange Server Component Management view (see Figure 26 on page 43).
2. Right-click the name of a connector, then select **Start ConnectorName**, **Pause ConnectorName**, or **Stop ConnectorName**.

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 ConnectorName	Starts the selected connector if it is paused or stopped. Connectors poll the application and connector controllers read the persistent queue. Flows are processed.
Pause ConnectorName	Pauses the selected connector if it is running or stopped. Connectors stop polling the application and connector controllers stop reading new requests, but events continue to be processed.
Stop ConnectorName	Stops the selected connector 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 ConnectorName	Shuts down the selected connector. The connector's process is stopped.
Boot ConnectorName	Restarts the selected connector. This action is available only if you have set the <code>OADAutoRestartAgent</code> property of the connector to <code>True</code> . See "Steps for setting automatic and remote restart for a connector" on page 48.

Steps for manually starting a connector

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

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

Note: To make a connector functional for the first time, you must configure it before you start the connector. See the adapter guide for the connectors you are using for configuration instructions.

Perform one of the following steps to start the connector:

- Click the desktop shortcut.
Start the connector by clicking the program shortcut created as part of the installation procedure.
- Select the connector's menu command in the IBM WebSphere submenu of the Windows **Start > Programs** 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 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 Express instance. Refer to the Connectors folder in System Manager for the correct *ConnectorName*.

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 7 to customize the startup of a connector.

Table 7. Connector startup parameters

Parameter	Function
-c <i>configFile</i>	Name of the configuration file to be used during startup. If the file name 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 Express (by default, <code>InterchangeSystem.cfg</code>).
-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 <code>.dll</code> file extension. The startup script specifies this option for all C++ connectors.
-f <i>pollFrequency</i>	Poll frequency is the number of milliseconds between polling actions. <ul style="list-style-type: none"> • To specify the number of milliseconds, provide a value for <i>pollFrequency</i>. • To cause the connector to poll only when you type the value <i>p</i> in the connector's Command Prompt window, specify the <code>-fkey</code> 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 <code>-fno</code>. <p>The value of this parameter overrides any repository definitions. You can specify either <code>-fkey</code> or <code>-fno</code>, but not both.</p>
-j	Specifies that the connector is written in Java. This parameter is optional if you specify <code>-l className</code> .
-l <i>className</i>	Specifies the name of the Java connector's global class, which is an extension of the connector base class. The startup script specifies this option for all Java connectors.
-n <i>connectorName</i>	Specifies the name of the connector to start.

Table 7. Connector startup parameters (continued)

Parameter	Function
-p <i>password</i>	Specifies the password that the connector uses to access InterChange Server Express.
-s <i>serverName</i>	Specifies the name of the server. This parameter is required. The name is case-sensitive.
-t	Turns on the connector property <code>SingleThreadAppCalls</code> . This property guarantees that all calls the connector framework makes to the application-specific connector code are with one event-triggered flow. The default value is <code>false</code> . Important: Do not change the value of this property from its shipped value. Each connector has the appropriate setting for its threading model. Specify this option only when starting a connector you created.
-x <i>connectorProps</i>	Passes application-specific connector properties to the connector. Use the format <code>prop_name=value</code> for each value you enter.

Shutting down a connector

The generic connector manager script calls the appropriate `start_connector.bat` script, which handles the actual connector management (including shutdown) for the connector. The InterChange Server Express product provides a `start_connector.bat` 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** check box is selected 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": the connector shuts down and does not restart until 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 the System View in System Monitor.

For instructions on enabling or disabling automatic restart, see "Steps for setting automatic and remote restart for a connector" on page 48.

Instructions for shutting down a connector depend on which tool you use to monitor the system:

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 for the connector that you want to shut down by selecting the check box to its left and click the Pause button in the upper-left corner of the view (see Figure 27 on page 45). Do this for each collaboration associated with the connector.

2. Select the check box for the connector that you want to shut down and click the Shutdown button in the upper-left corner of the view (see Figure 27 on page 45).

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

1. Expand the Collaboration Objects folder in the InterChange Server Component Management view (see Figure 30 on page 53).
2. Right-click the collaboration object associated with the connector and click **Pause *CollaborationName***. Do this for each collaboration associated with the connector.

The collaboration object icon changes to an icon with two vertical lines.

3. From the expanded Connectors folder in the InterChange Server Component Management view, right-click the connector and click **Shut Down *ConnectorName***.

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 use the Ctrl+C key sequence, or if you use “q,” or other manual methods to perform the shutdown, and the Object Activation Daemon (OAD) is enabled, OAD immediately restarts the connector.

Restarting a connector

The procedures described in this section restart the connector after you have used the Shut Down Connector action in either System Monitor or in System Manager in the InterChange Server Component Management view. 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”).

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

1. From the System Overview view (see Figure 25 on page 43), select the check box to the left of the connector that you want to restart.
2. Click the Restart Agent button in the upper-left corner of the view (see Figure 27 on page 45).

Steps for using System Manager to restart a connector: Perform the following steps in System Manager to restart a connector:

1. Expand the Connectors folder in the InterChange System Component Management view of System Manager, (see Figure 26 on page 43).
2. Right-click the connector that you want to restart, then click **Boot *ConnectorName***.

Steps for setting automatic and remote restart for a connector

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

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

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. (To manually start the MQ-trigger Monitor: use the runmqtrm command, or the runmqtmc command if the MQ client is started on that machine.)
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 8.

Table 8. 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 8 on page 49.
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 using OAD as a Windows service

If you are configuring a connector for automatic restart or remote startup using OAD, do not configure the connector to run as a automatic Windows service. Instead, install WebSphere MQ Trigger Monitor to run as a Windows service. When the system starts, OAD automatically starts as a Windows service; when InterChange Server Express restarts, it starts the connector through the OAD.

Important: By default, connectors are installed as Windows services which need to be manually started. The remote restart feature does not function if the connector’s Windows service is changed from manual to automatic.

Perform the following steps to install WebSphere MQ Trigger Monitor as a Windows service:

1. Select **Start > Programs > IBM WebSphere MQ > WebSphere MQ Services**.
2. Right-click the queue manager that WebSphere Business Integration Server Express is using, and click **New > Trigger Monitor**. The Create Trigger Monitor Service dialog box appears.
3. Click the **Parameters** tab. In the queue name field, type the string “INITIATION.QUEUE”. In the **Description** column, Trigger Monitor appears as one of the services for that queue manager.
4. Run the Windows Service Setup tool. Select **InterChange Server Express** in the **Service Component** list
5. In the **Service Dependencies** field, type the name for your OAD Windows service.

After completing these steps, you can use the Services tool in Windows to start and stop the daemon.

Steps for configuring flow control for connectors

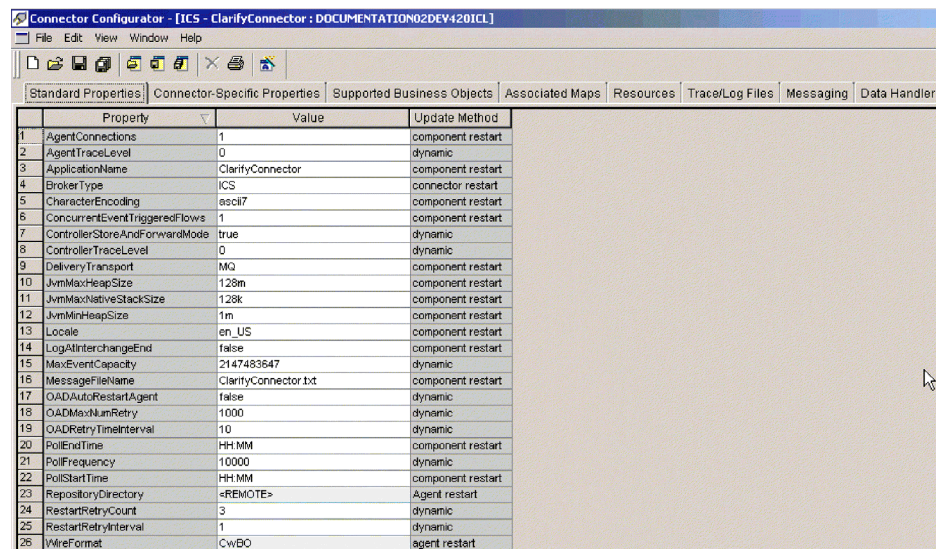
Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control system-wide, see “Steps for configuring system-wide flow control” on page 40. 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 Express to be rebooted. System-wide configuration changes for flow control require InterChange Server Express 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 in System Manager in the InterChange Server Component Management view. 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 16. For more information on viewing the flow control from the InterChange Server Component Management view, see “Steps for viewing collaboration object statistics” on page 28 or “Steps for viewing connector statistics” on page 30.

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

1. In System Manager, navigate to the connector for which you want to configure flow control, then double-click that connector. Connector Configurator Express opens (see Figure 28).



Property	Value	Update Method
1 AgentConnections	1	component restart
2 AgentTraceLevel	0	dynamic
3 ApplicationName	ClarifyConnector	component restart
4 BrokerType	ICS	connector restart
5 CharacterEncoding	escl7	component restart
6 ConcurrentEventTriggeredFlows	1	component restart
7 ControllerStoreAndForwardMode	true	dynamic
8 ControllerTraceLevel	0	dynamic
9 DeliveryTransport	MQ	component restart
10 JvmMaxHeapSize	128m	component restart
11 JvmMaxNativeStackSize	128k	component restart
12 JvmMinHeapSize	1m	component restart
13 Locale	en_US	component restart
14 LogAllInterchangeEnd	false	component restart
15 MaxEventCapacity	2147483647	dynamic
16 MessageFileName	ClarifyConnector.txt	component restart
17 OADAutoRestartAgent	false	dynamic
18 OADMaxNumRetry	1000	dynamic
19 OADRetryTimeInterval	10	dynamic
20 PollEndTime	HH:MM	component restart
21 PollFrequency	10000	dynamic
22 PollStartTime	HH:MM	component restart
23 RepositoryDirectory	<REMOTE>	Agent restart
24 RestartRetryCount	3	dynamic
25 RestartRetryInterval	1	dynamic
26 WireFormat	CwBO	agent restart

Figure 28. Connector Configurator Express, Standard Properties tab

2. In the Standard Properties tab, click in the **Value** cell of the MaxEventCapacity property.

3. Change the value to the maximum number of events that you want queued for a connector. The valid range of values for this property is from 1 to 2147483647.
4. Click **File > Save > to Project**. The following message appears in the bottom section of Connector Configurator Express: Connector '*ConnectorName*' is saved successfully.

Administering collaboration objects

Administering 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 from System Manager in the InterChange Server Component Management view.

This section covers the following topics:

“Viewing collaboration object states” on page 52

“Starting, stopping, and pausing collaboration objects” on page 54

“Configuring collaboration object run-time properties” on page 55

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 in System Manager. To log on to System Monitor, follow the instructions in “Steps for logging on to System Monitor” on page 13.. To use System Manager, follow the instructions in “Steps for connecting to an InterChange Server Express instance” on page 25..

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

Steps for using System Monitor to view collaboration object states

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

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

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 16 for instructions.

2. Click the triangle next to the name of the server to display a list of components on the system. All collaboration objects are listed along with their status, start time, and total up time (see Figure 29 on page 53).

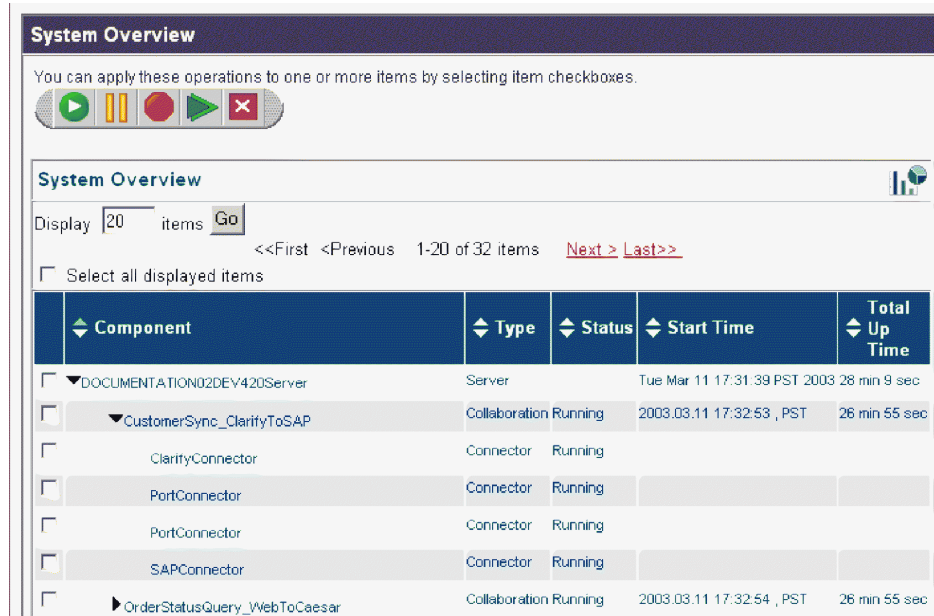


Figure 29. System Monitor, System Overview displaying collaboration object status

Steps for using System Manager to view collaboration object states

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

1. Open System Manager and connect to the InterChange Server Express instance that contains the collaboration object that you want to view (see “Steps for connecting to an InterChange Server Express instance” on page 25).
2. In the InterChange Server Component Management view, expand the InterChange Server Express 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 30).

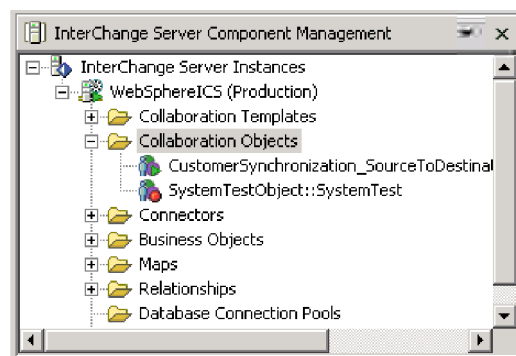


Figure 30. Collaboration Objects folder in the InterChange Server Component Management view

Table 9 describes the collaboration object states that are viewable from System Monitor and the Collaboration Objects folder in the InterChange Server Component Management view:

Table 9. Collaboration object states

Collaboration object state	Description
Start	A started collaboration subscribes to its triggering business objects and processes them as they arrive. If you stop and then restart InterChange Server Express, collaboration objects in the Start state automatically start running when InterChange Server Express restarts.
Pause	A paused collaboration cannot receive new flow initiators. It 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 resumes. To resume collaboration execution, click Start in System Monitor or in the Collaboration Object menu in the InterChange Server Component Management view.
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 when 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 run a collaboration object for the first time, you must first configure it then start it. See “Configuring collaboration object run-time properties” on page 55 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 in System Monitor to start, stop, and pause collaboration objects:

1. In the System Overview view (see Figure 25 on page 43), select a collaboration object by selecting the check box to its left.

2. Click the Start, Pause, or Stop icon from the icon group in the upper-left corner of the view (see Figure 27 on page 45).

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

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

1. Go to the Collaboration Objects folder of the InterChange Server Component Management view.
2. Right-click a collaboration object, then select **Start CollaborationName**, **Pause CollaborationName**, or **Stop CollaborationName**.

Configuring collaboration object run-time properties

This section describes system administration tasks in a production environment and contains the following procedures:

“Steps for setting collaboration object general properties” on page 55

“Steps for configuring collaboration objects to process concurrent event-triggered flows” on page 57

“Steps for configuring flow control for collaboration objects” on page 57

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:

1. In System Manager, expand the Collaboration Objects folder in the InterChange Server Component Management view.
2. Right-click a collaboration object and click **Properties**.
3. In the Properties dialog box, click the **Collaboration General Properties** tab. See Figure 31 on page 56.

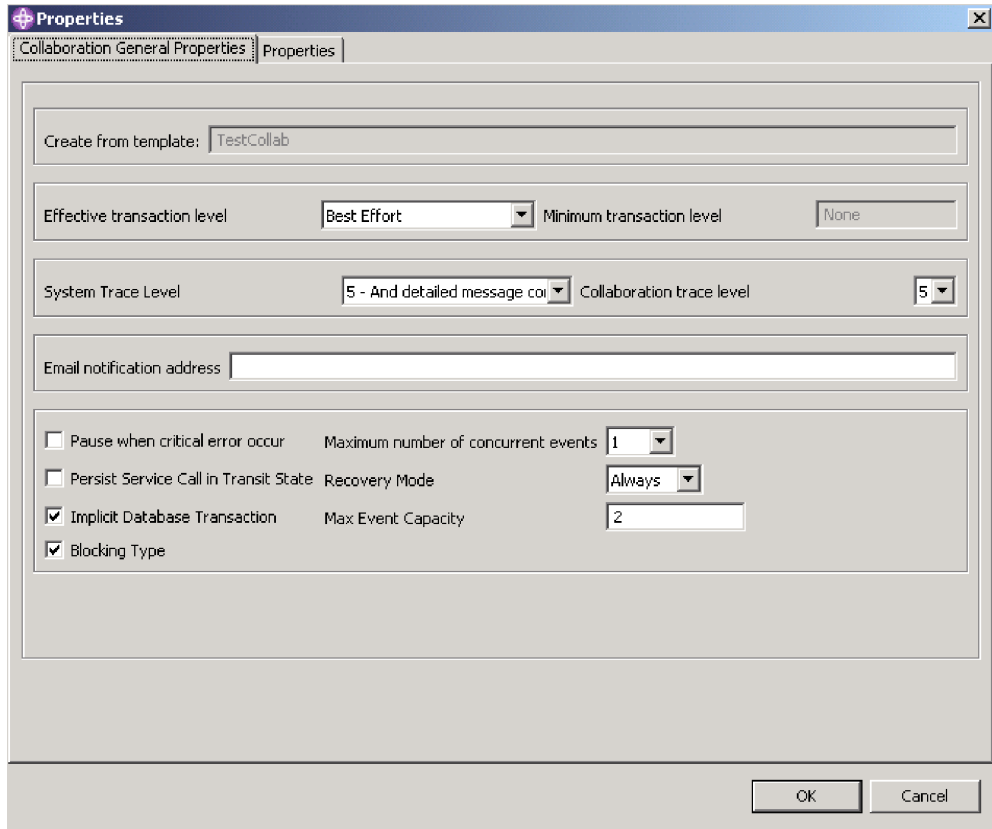


Figure 31. Properties dialog box, Collaboration General Properties tab

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 set the following:

- **Effective transaction level**
See the *System Implementation Guide*.
- **System trace level**
See “Configuring logging and tracing” on page 110.
- **Collaboration trace level**
See “Steps for configuring the collaboration object trace level” on page 114.
- **E-mail notification address**
See “Steps for configuring e-mail notification at the collaboration object level” on page 109.
- **Pause when critical error occurs** check box
See “Critical errors” on page 126.
- **Persist Service Call in Transit State** check box
See Chapter 3, “Troubleshooting the system,” on page 99, and see the *Collaboration Development Guide*.
- **Implicit Database Transaction** check box
See the *System Implementation Guide*.
- **Blocking Type** check box
See the *System Implementation Guide*.

- **Maximum number of concurrent events**
See “Steps for configuring collaboration objects to process concurrent event-triggered flows” on page 57.
- **Recovery Mode**
See “Steps for deferred recovery of collaboration events” on page 124.
- **Max Event Capacity**
See “Steps for configuring flow control for collaboration objects.”

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 and click **Properties**. The Properties dialog box appears (see Figure 31 on page 56).
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 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 40. 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 Express to be rebooted. System-wide configuration changes for flow control require InterChange Server Express 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 in System Manager from the InterChange Server Component Management view. 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 16. For more information on viewing the flow control in System Manager from the InterChange Server Component Management view, see “Steps for viewing collaboration object statistics” on page 28 or “Steps for viewing connector statistics” on page 30.

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

1. In System Manager, expand the Collaboration Objects folder in the InterChange Server Component Management view.
2. Right-click the collaboration object for which you want to create flow control and click **Properties**. The Properties dialog box appears (see Figure 31 on page 56).
3. 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, inclusive.
4. 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 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. Click 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 is selected, the value can be edited.

Note: The name of the service call timeout configuration property may be something like, "CreateTimeout" or "RetrieveTimeout," but because there is no naming convention for this property, you may have to contact the collaboration developer if the name of the service call timeout configuration property is not immediately apparent.

3. Edit the value so that it represents 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 wait time is equal to infinity. If it contains non-numerical values, it triggers a collaboration run-time exception.

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

Administering maps

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

This section covers the following topics:

"Viewing map states" on page 59

"Starting and stopping maps" on page 60

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 in System Manager. To log on to System Monitor, follow the instructions in “Steps for logging on to System Monitor” on page 13. To use the InterChange Server Component Management view of System Manager, follow the instructions in “Steps for connecting to an InterChange Server Express instance” on page 25..

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 in System Monitor to view map states:

1. Open the System Overview view.
2. In the **Views** list, click the **Maps and Relationships** link. The Map Status and Relationship Status monitors appear (see Figure 11 on page 15).

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 16 for instructions.

Steps for viewing map states in System Manager

Perform the following steps in System Manager to view map states:

1. Connect to the InterChange Server Express instance that contains the map that you want to view. See “Steps for connecting to an InterChange Server Express instance” on page 25 for instructions on connecting to an InterChange Server Express instance.
2. In the InterChange Server Component Management view, expand the InterChange Server Express instance, and then expand the Maps folder. The maps appear under the expanded Maps folder with different colored lights to indicate their different states (see Figure 32).

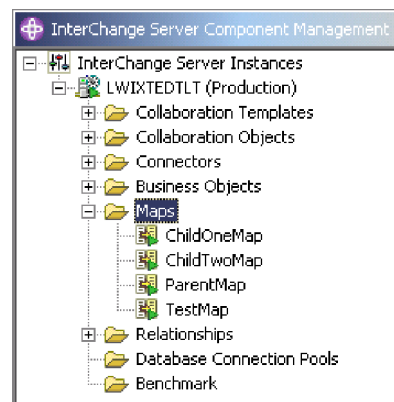


Figure 32. Maps folder in the InterChange Server Component Management view

Table 10 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 10. Map States

Map State (color)	Description
Active (green)	Map is ready to run and available for use.
Inactive (red)	Map is not ready to run or available for use.
Unknown and does not display	Map has not been compiled or the map class is missing. The map is not ready to run or available for use. Compiling the map in Map Designer Express puts the map in active state, when 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 InterChange Server Express 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, refer to the *Map Development Guide*.

This section describes how to start and stop maps. For information about additional map tasks, 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 Express 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 InterChange Server Express 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 the following way:

Perform the following steps in System Monitor to start and stop maps:

1. If the System Overview view is displayed, click the **Maps and Relationships** link under Views in the left pane. The Map Status and Relationship Status monitors appear (see Figure 11 on page 15).
2. Select the Start or Stop icon in the upper-left corner of the view (see Figure 27 on page 45).

Perform the following steps in System Manager to start and stop maps:

1. Expand the Maps folder in the InterChange Server Component Management view (see Figure 32 on page 59).
2. Right-click a map, then select either **Start MapName** or **Stop MapName**.

Administering relationships

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

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 in System Manager. To log on to System Monitor, follow the instructions in “Steps for logging on to System Monitor” on page 13. To use the InterChange Server Component Management view of System Manager, follow the instructions in “Steps for connecting to an InterChange Server Express instance” on page 25..

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

Steps for using System Monitor to view relationship states

Perform the following steps in System Monitor to view relationship states:

1. Open the System Overview view.
2. Under **Views**, click the **Maps and Relationships** link. The Map Status and Relationship Status monitors appear (see Figure 11 on page 15).

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 16 for instructions.

Steps for using System Manager to view relationship states

Perform the following steps in System Manager to view relationship states:

1. Connect to the InterChange Server Express instance that contains the relationship you want to view. See “Steps for connecting to an InterChange Server Express instance” on page 25 for instructions on connecting to an InterChange Server Express instance.
2. In the InterChange Server Component Management view, expand the InterChange Server Express instance, then expand the Relationships folder and expand either the Dynamic or Static folder. The relationships appear under either folder and have different icons to indicate their different states (see Figure 33).

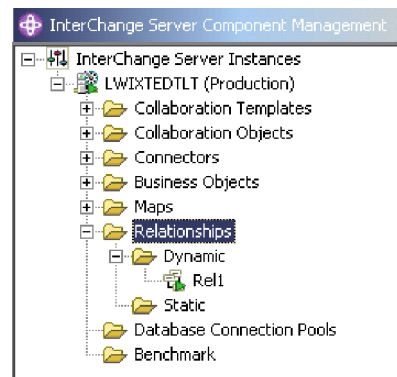


Figure 33. Relationships folder in the InterChange Server Component Management view

Table 11 lists the relationship states represented by the display color and describes what actions can be performed during that state.

Table 11. Relationship States

Relationship state (color)	Description
Active (green)	Relationship is ready to run and available for use. 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. When the relationship is stopped, it becomes inactive -- all current jobs in the queue are completed and no new jobs are accepted. To modify a relationship definition, it must be in this state.
Unknown (gray)	Relationship does not have a compatible run-time schema. To create a compatible run-time schema, from Relationship Designer Express, save the relationship with the Create run-time schema option selected. The state changes to Inactive, at which point the relationship can 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 Express. For more information about Relationship Designer Express, see the *Map Development Guide*.

When you expand the Relationships folder in the InterChange Server Component Management view, 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 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 62

“Steps for starting and stopping relationships”

“Relationship table caching” on page 63

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, modify it, and then restart it. See Appendix B, “Requirements for restarting IBM WebSphere InterChange Server Express system components,” on page 157 to find out what changes require relationships to be restarted.

Steps for starting and stopping relationships

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

1. From the System View window (see Figure 11 on page 15), select **Show Maps and Relationships** from the **View** menu. The **Maps** and **Relationships** columns appear next to **Collaborations** and **Connectors**.
2. Right-click a relationship, then select either **Start RelationshipName** or **Stop RelationshipName**.

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

1. Expand the Dynamic or Static folder in the InterChange Server Component Management view (see Figure 33 on page 61).
2. Right-click the name of a relationship, and click **Start RelationshipName** or **Stop RelationshipName**.

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 so is a good candidate for table caching. 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 labeled Static in the Relationships folder.

Note: For information on how to design a static relationship so that its tables are 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. Select **Cached** 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. Select **Cached** 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 tell InterChange Server to 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. Select **Reload** to reload the static relationship's tables.

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, select the **Reload** option.

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

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

```
RELATIONSHIP.CACHING=5
```

InterChange Server Express routes these messages to the trace file (if one is configured). By default, InterChange Server Express 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 Express. 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 start. The data is stored in the relationship tables specified in the relationship definition. Relationship Manager provides a graphical interface to interact with the relationship tables regardless of the database vendor.

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

Figure 34 on page 65 shows an example in Relationship Manager of a relationship tree for an identity relationship.

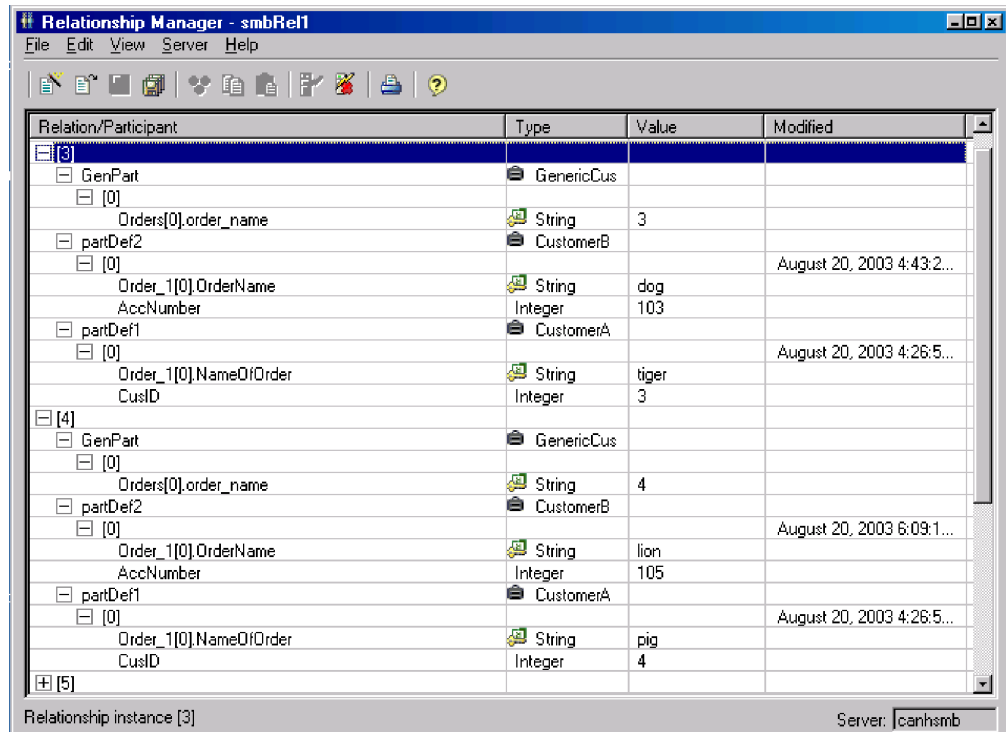


Figure 34. Relationship Manager, relationship tree

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.

Steps for starting Relationship Manager

Perform one of the following steps to start Relationship Manager:

- Select **Start > Programs > IBM WebSphere Business Integration Express > Toolset Express > Administrative > Relationship Manager**
- In Relationship Designer Express, select a relationship definition and then select **Tools > Relationship Manager**

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

Figure 35 shows Relationship Manager in a disconnected state.




Figure 35. Relationship Manager

Connecting to and disconnecting from a server

You must connect Relationship Manager to InterChange Server Express 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 Relationship Manager to InterChange Server Express”
- “Steps for disconnecting Relationship Manager from InterChange Server Express” on page 67

Steps for connecting Relationship Manager to InterChange Server Express:

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

1. Click **Server > Connect**. The Connect to InterChange Server dialog box appears (see Figure 36 on page 67).



Figure 36. Connect to InterChange Server Express

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

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

- Select a cached server name from the list.
 - Click the browse button to browse for the InterChange Server Express instance on the network and click **OK**.
3. Type the user name to interact with the InterChange Server Express 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 Express instance in System Manager then enable 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 68.

7. Click **Connect**.

If you connect to InterChange Server Express in Relationship Manager and specify a relationship to open as described in step 6, then Relationship Manager displays the Retrieve Relationship Instances window, described in “Steps for retrieving relationship instances” on page 68.

Steps for disconnecting Relationship Manager from InterChange Server Express:

Perform the following step to disconnect Relationship Manager from InterChange Server Express:

- Click **Server > Disconnect**.

Working with relationships in Relationship Manager

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

“Steps for opening a relationship”

“Steps for retrieving relationship instances”

“Steps for creating relationship instances” on page 71

“Steps for deleting relationship instances” on page 72

“Deactivating and activating participants” on page 72

“Steps for copying participants” on page 73

“Steps for loading and unloading business object files” on page 74

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. In Relationship Manager, click **File > Open**.
2. In the Open Relationship dialog box, select the name of the relationship that you want to open.

Figure 37 shows the Open Relationship dialog box.

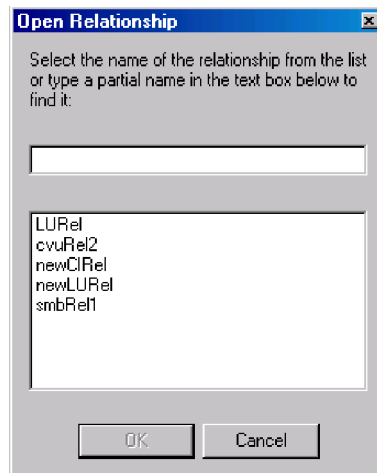


Figure 37. Opening Relationship dialog box

3. Click **OK**.

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

Steps for retrieving relationship instances

Perform the following steps to retrieve relationship instances or to return a count of how many instances there are for a relationship:

1. In Relationship Manager, click **File > Retrieve**.
The Retrieve Relationship Instances dialog box opens (see Figure 38).

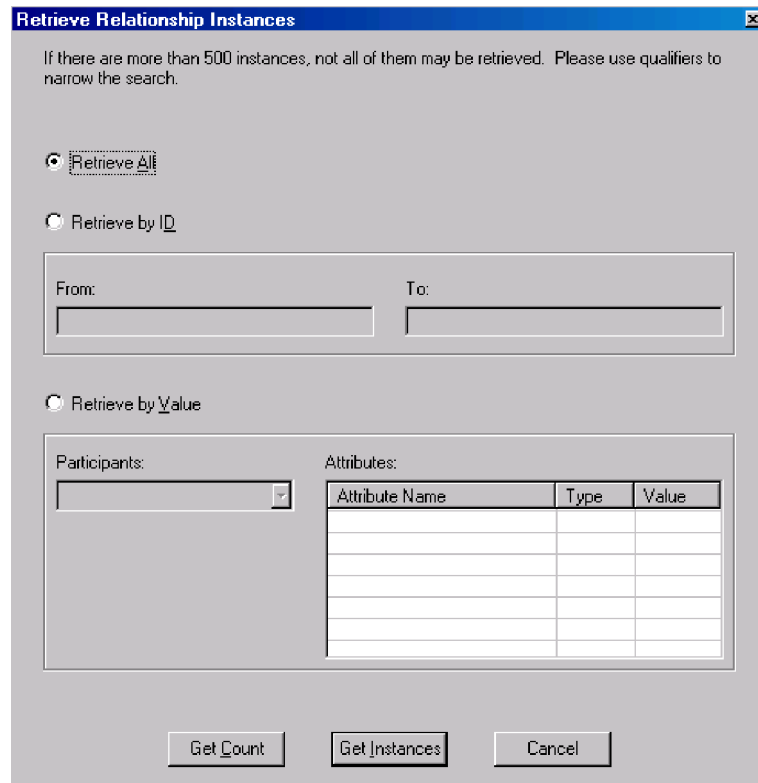


Figure 38. Retrieving Relationship Instances dialog box

2. 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 instances by relationship ID” on page 70.
 - Retrieve a relationship instance that contains a participant of a particular value you specify, as described in “Steps for retrieving relationship instances by participant data” on page 70.
 - Return a count of the number of instances for the relationship, as described in “Steps for returning a count of relationship instances” on page 70.

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

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 instances 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 specified.

Steps for retrieving relationship instances 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** list.

For identity relationships, participant names are followed by the business object definition with which the participant is associated.

For lookup relationships, participant names are followed by the word "Data".

3. Type one of the types of values listed in Table 12 in the **Value** column in the **Attributes** pane.

Table 12. 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 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 69, “Steps for retrieving instances by relationship ID” on page 70, or “Steps for retrieving relationship instances by participant data” on page 70
2. Click **Get Count** (instead of **Get Instances**).

Steps for creating relationship instances

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

1. Create the new relationship instance by doing one of the following:
 - Click **File > New Instances**.
 - 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 (???. After you save the relationship instance or any of its participants, InterChange Server Express 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 click **Add Participant**.
 - Click **Add Participant** on the toolbar.
 - c. Expand the new participant instance by clicking the plus (+) sign next to it.
 - d. Select the new participant instance.
 - e. Click the **Value** column for the participant instance, then type a value into the cell.

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

4. Do any of the tasks listed in Table 13 on page 72.

Table 13. Tasks for participant data

Task	Action
Save the participant instance.	To save the new participant instance, right-click the participant instance and click Save Participant . 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: After the participant data has been saved, it cannot be changed. To change its data, the participant must be deleted and another created.
Add more participant instances.	Repeat repeat step 3 on page 71 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 the participant instance and selecting Delete Participant . 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 72). A deactivated participant retains its instance ID and its values.
Save the relationship instance.	Save the relationship instance by performing one of the following tasks: <ul style="list-style-type: none"> • Click File > Save. • Right-click the relationship instance and select Save Relationship. InterChange Server Express 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.	Click File > Save All . InterChange Server Express 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 that you want to delete.
2. Do one of the following in Relationship Manager:
 - Click **File > Delete**.
 - Right-click the relationship instance and click **Delete Relationship Instance**.

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 step to deactivate a participant instance:

- Right-click the participant instance you want to deactivate and click **Deactivate Participant**.
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**.

The Deactivated Participants window appears as shown in Figure 39.

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 39. Deactivated Participants window

2. Select the relationship instance that contains the deactivated participant you want to activate from the list.
3. Expand the relationship instance until the deactivated participant instances appear in the list.
4. Right-click the participant instance that you want to reactivate and select **Activate**.
5. Click **Edit > Refresh**.

The activated participant instance appears 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 click **Add Participant**.
2. Right-click the participant instance that you want to copy and click **Copy Participant**.
3. Right-click the newly created participant instance and click **Paste Participant**.

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 click **Load Participant with Business Object**.

The Load Participant with Business Object dialog box displays the business object associated with that participant instance, as shown in Figure 40.




Figure 40. Load Participants with Business Objects dialog box

2. Click **Load**.
3. Navigate to and open the business object file that 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 in Relationship Manager

An important feature of Relationship Manager is the 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 75

“Steps for filtering the displayed participants” on page 76

“Steps for cleaning up participants” on page 77

“Steps for printing relationship data” on page 78

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.

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

1. Select a participant instance in Relationship Manager.
2. Click **Edit > Find Instances by Business Object**.

The Find Instances by Business Object window appears (see Figure 41).




Figure 41. Find Instances by Business Object dialog box

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 to navigate to and select the instance.

Perform the following steps to search for instances by data:

1. Select a participant instance in Relationship Manager.
2. Click **Edit > Find Instances by Data**.

Relationship Manager displays the Find Instances by Data dialog box, as shown in Figure 42.




Figure 42. Find Instances by Data dialog box

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 to navigate to and select 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. Click **View > Filter**.
The Filter dialog box appears, as shown in Figure 43.


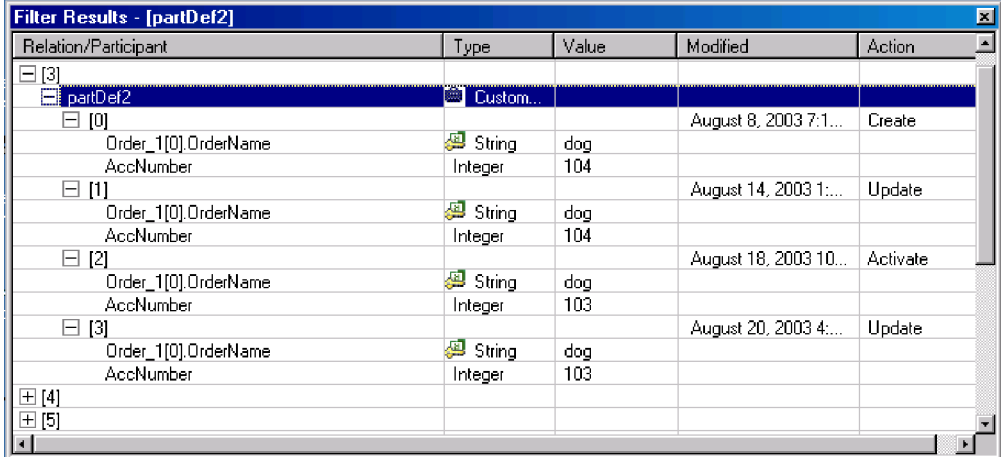


Figure 43. Filter dialog box

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.
 - In the months list, select a month. You can type the first letter of a month to scroll to a month in the list.
 - Click the up and down arrows to increase or decrease the date. You also can use the arrow keys on the keyboard.
4. Select the **Include inactive participants as well** check box if you want to include inactive participants in the results.

5. Click **OK**.

Relationship Manager displays the history of activity for the filtered interval in the Filter Results dialog box. The dialog box includes inactive participants if **Include inactive participants as well** is selected. Figure 44 shows the Filter Results dialog box.



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 44. Filter Results dialog box

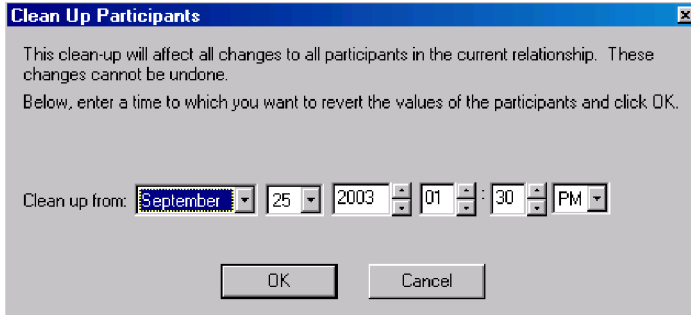
Steps for cleaning up participants

You can remove incorrect participant information caused by inconsistent or corrupt data from the source application or generic object.

Perform the following steps to clean up participants in System Manager:

1. Select a participant in Relationship Manager.
2. Click **Edit > Clean Up Participants**.

The Clean Up Participants dialog box appears, as shown in Figure 45.



Clean Up Participants

This clean-up will affect all changes to all participants in the current relationship. These changes cannot be undone.

Below, enter a time to which you want to revert the values of the participants and click OK.

Clean up from: September 25 2003 01 : 30 PM

OK Cancel

Figure 45. Clean Up Participants dialog box

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.
 - In the months list, select a month. You can type the first letter of a month to scroll to a month in the list.
 - Click the up and down arrows to increase or decrease the date. You also can use the arrow keys on the keyboard.
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 print command in 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, designate those instances by selecting them.
3. Print relationship run-time data in any of the following ways:
 - Click **File > Print**.
 - Use the keyboard shortcut **Ctrl+P**.
 - Click the **Print** icon in the toolbar.
4. The Print Relationship Instances dialog box appears. Select either all instances or selected instances, then click **OK**.

Scheduling jobs in the system

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 Express 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 78

"Steps for creating schedules" on page 81

"Steps for modifying schedules" on page 82

"Steps for deleting schedules" on page 83

"Steps for displaying schedules" on page 83

"Steps for enabling or disabling schedules" on page 83

Overview of scheduling jobs

You schedule jobs in System Manager in the Schedule window (see Figure 46). From the Schedule window, you can create, modify, and delete scheduled items. You can see a list of all the schedules that are defined for components, or selectively view schedules based on your requirements. You can also enable or disable all schedules on the server.

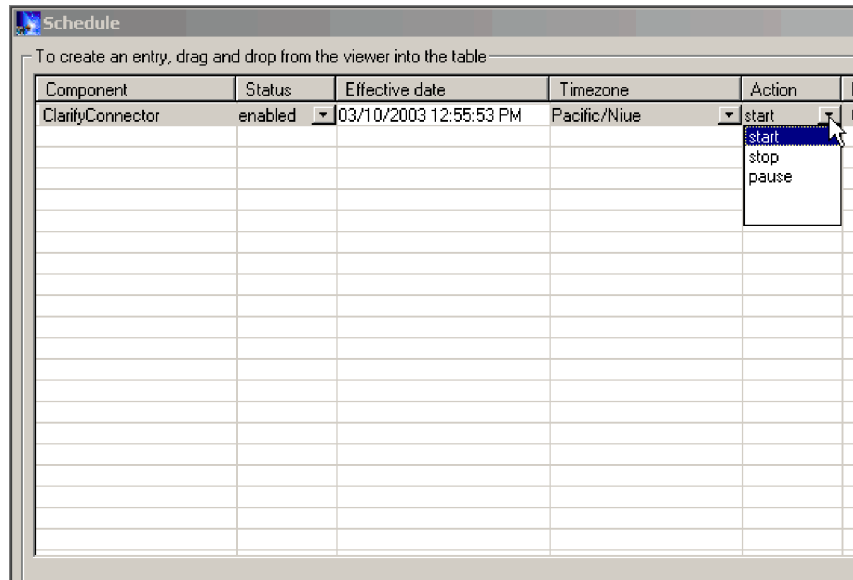


Figure 46. System Manager Schedule window

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. After a schedule is set, you can enable or disable its use.

The Schedule window allows you to determine the following items:

- Component The name of the connector or collaboration being scheduled.
- 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 (PST).
- 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.
- Comments Text field that contains comments you enter about the schedule.

If you want 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 Express 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., depending on the PollFrequency setting. When you pause a connector, it stops polling until started again, but is still able to handle service call requests from InterChange Server Express. 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 52. 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 47, each collaboration object gets a dedicated time period when the connector is processing only that collaboration object's events.




Figure 47. Dedicated processing

Multiple collaboration objects can subscribe to the same business object. In that case, the object is sent to InterChange Server Express, 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 that the scheduler cannot change. For example, 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 Schedules folder in System Manager, then clicking **Edit Components' schedule**. The Schedule window appears, as shown in Figure 46.

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, select the date and time when the scheduled item occurs.

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 (available from the Edit menu).

- c. In the **Timezone** field, select 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) and InterChange Server Express is located in Japan. InterChange Server Express 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 comments in the **Comments** field. A maximum of 255 characters is allowed.
4. If you want this schedule to be ongoing, select the **Recurrence** check box and enter information about the next occurrence of the action. Select 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:
 - **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 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 Express automatically handles changes between standard and daylight savings time for recurring events.

5. Select the **Show Dependencies** check box to display schedules for a collaboration object's bound connectors and collaborations
6. Select the **Show Expired** check box to display schedules that have already processed and whose time to run has expired. Only non-recurring schedules expire.
7. Click **OK** or **Apply** to create the scheduled item, which is effective after InterChange Server Express receives the information.

When InterChange Server Express and components are geographically distant, there can be a slight delay. If you need to change the state of a component immediately, use System Monitor, rather than the scheduler, to start, stop, or pause a component.

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 78 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. In System Manager, right-click the collaboration or connector, and click **Edit Components' schedule**. The Schedule window appears (see Figure 46).
2. Edit any field in the Schedule 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 appear 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. In System Manager, right-click the collaboration or connector, and click **Edit Components' schedule**. The Schedule window appears (see Figure 46).
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 one of the following objects for displaying schedules:
 - The collaboration or connector icon in 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.
 - The Collaboration Objects or Connectors folder in System Manager. The Schedule window displays schedules for all objects in that folder.
 - InterChange Server Express in System Manager. The Schedule window displays all schedules defined for that server, with times shown for the server's time zone.
2. Click **Component > Schedule** from the main window to display all schedules in the system.
3. Click any of the column headings in the Schedule window 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 83..
2. Enable or disable the schedule by doing one of the following in the Schedule window:
 - To enable or disable all schedules, click either **Enable All** or **Disable All**.
 - To enable or disable a single schedule, click the down arrow in the **Status** column and click the **Enable** or **Disable** option.
3. Click **Apply** to complete this task.
4. Click **OK** to exit.

Backing up system components

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

This section covers the following topics:

“Backup schedule planning” on page 84

“Component backups” on page 85

Backup schedule planning

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

Within the InterChange Server Express 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, back up static configuration data stored in the InterChange Server Express repository only when it has changed or before planned reinstallations or upgrades to the system. Following is a partial list of static data in the InterChange Server Express system:
 - InterChange Server Express 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. Because relationship instance data is maintained continuously, regularly back up this data and 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 InterChange Server Express system:

- InterChange Server Express cross-reference database
- InterChange Server Express relationship tables
- InterChange Server Express WIP (event management) and transaction tables
- WebSphere MQ queue data
- InterChange Server Express 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. InterChange Server Express 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 InterChange Server Express environment require different backup procedures. The following topics are described in this section:

“Relationship table backups” on page 85

“Repository backups” on page 85

“System installation file backups” on page 86

“Collaboration class file backups” on page 86

“Archive table backups” on page 86

Attention: When backing up IBM WebSphere InterChange Server Express 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, InterChange Server Express 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 InterChange Server Express system is in a quiescent state when backing up the relationship tables. For more information on bringing the system to a quiescent state, see “Shutting down InterChange Server Express” on page 35.

It is recommended that you mirror the relationship database log to assist in recovery. If hardware/software cost is not a consideration, you can also mirror the relationship run-time data.

The set of relationship tables for one relationship are closely associated, so 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 86. Back up the repository whenever it is modified and before and after performing a reinstallation or an InterChange Server Express software upgrade. The InterChange Server Express 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 you mirror the Repository, Event Management, and Transaction database logs to assist in recovery.

Unpartitioned (single) database backups: If your InterChange Server Express databases are not partitioned, meaning they are contained in a single database, they should not be part of your normal database backup routine. The InterChange Server Express databases contain transient data whose recovery can cause inconsistencies in the system. Instead, back up the objects in the InterChange Server Express repository by using the `repos_copy` utility.

System installation file backups

Back up the system installation files 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-InterChange Server Express 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 InterChange Server Express 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 InterChange Server Express system. These events can be used to "resynchronize" the application and the InterChange Server Express cross-reference tables. For more information, see "Maintaining event archives" on page 122.

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. Refer to the *WebSphere Business Integration Server Express Installation Guide* for more information. `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 in an MS-DOS command prompt window. 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` does not work.

Important: When `repos_copy` deploys components to the repository, it deploys them to the repository *only*. It does *not* deploy them to any in-memory tables of business object definitions. For example, 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 stop and restart InterChange Server Express 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 87
- “`Repos_copy` usage scenarios” on page 91
- “Locale for `repos_copy` files” on page 97

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

Repos_copy syntax

Table 14 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, at least one of which is 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.

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

Table 14. `Repos_copy` command options

Option	Description
<code>-ai</code>	Ignore and do not load duplicate objects (business objects, maps, relationships, collaboration templates and objects, and connectors) that are found when deploying a package.
<code>-ar</code>	Replace duplicate objects (business objects, maps, relationships, collaboration templates and objects, and connectors) that are found when deploying a package.

Table 14. *Repos_copy* command options (continued)

Option	Description
-arp	Replace duplicate components. This is an interactive version of the -ar option. If the components in the package being deployed already exist in the repository, you are prompted to ignore or replace the component.
-d	Delete the components in the repository, except the state data. Use this option to delete all of the components from the repository.
-deep	Use with the -e option to include all the dependent components. If you omit the -deep option, only the component that is specified with the -e option is included.
-dfo <i>EntityType:Entity</i> [+ <i>EntityType:Entity2</i>]	Forcefully deletes 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. See also the -do option.
-do <i>EntityType:Entity</i> [+ <i>EntityType:Entity2</i>]	Specify the entities to be deleted from the repository. See Table 15 on page 91 for the list of entity types and keywords. If the object has no referents—other components that depend on it—then the deletion takes place. If the object has referents, then the deletion fails and a message is displayed. The behavior is the same in both design mode and production mode. For more information about starting the server in design mode or production mode, see the <i>System Implementation Guide</i> .
-e <i>EntityType:Entity1</i> [+ <i>EntityType:Entity2...</i>]	<p>Export 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 15 on page 91.</p> <p>Follow the <i>EntityType</i> keyword with a colon (:) and the name of the entity. Use the “+” to specify more than one entity. When combined with the -o option, the -e option unloads the data to an output file.</p>

Table 14. *Repos_copy* command options (continued)

Option	Description
<i>-fentityFile</i>	<p>Similar to the <i>-e</i> option, but take the names of the entities from the file <i>EntityFile</i>. The file should contain references to the entities, with the following conditions:</p> <ul style="list-style-type: none"> • The entity names must follow after the proper entity type keyword. The entity types and their keywords are listed in Table 15 on page 91. • A colon must separate the entity type from the entity name. • There must be a new line separating each entity reference. <p>When combined with the <i>-o</i> option, this option exports the components to a package.</p>
<i>-ifilename</i>	<p>Deploy 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 <i>.jar</i> file containing objects in XML format or an <i>.in</i> file containing objects in a text format.</p> <p>The <i>.jar</i> files created by <i>repos_copy</i> or System Manager have a particular structure which must be maintained for any subsequent imports of such a file to be successful. You should not, therefore, ever modify an input file manually.</p>
<i>-k</i>	<p>Override the default behavior of <i>repos_copy</i> when it finds a Mercator map in the package file it is loading. By default, <i>repos_copy</i> exits if it encounters a Mercator map. If you use the <i>-k</i> option, <i>repos_copy</i> skips over any Mercator maps in the package file and proceeds with the deployment process.</p>
<i>-mode</i>	<p>Return the mode of the server. For more information about InterChange Server Express modes, see the <i>Implementation Guide for WebSphere InterChange Server</i>.</p>
<i>-ncencoding</i>	<p>See the Java documentation about the String class for a list of valid character encodings.</p> <p>When deploying a repository created with WICS version 4.1.1, specify the character encoding used in the 4.1.1 environment.</p>

Table 14. *Repos_copy* command options (continued)

Option	Description
<i>-outfilename</i>	Export the components in the repository to the specified package file. You must specify the name of the package file. If the file already exists then <i>repos_copy</i> prompts you to overwrite it or not. The output file is in <i>.jar</i> format, and contains the component definitions in XML format, and <i>.java</i> source files for components that have them. This option cannot be combined with the <i>-i</i> or <i>-d</i> options, nor can it export components in text format as it did in previous releases. <i>Repos_copy</i> does not append the <i>.jar</i> extension, so you must specify it when specifying the name of the output file.
<i>-ppassword</i>	Specifies the password for the user name supplied with the <i>-u</i> option. The password case-sensitive. If you do not specify this option then <i>repos_copy</i> prompts you for the password.
<i>-r*</i>	Import relationship definitions and not create the run-time schemas for any of them. See also the <i>-r</i> option.
<i>-rrelationshipName1[:relationshipName2]</i>	Load the named relationship definition(s) into the repository without creating its run-time schema.
<i>-sserverName</i>	Specify the name of the InterChange Server Express instance with which <i>repos_copy</i> should interface. The name is case-sensitive. If the server name is not specified, the tool prompts for a server name.
<i>-summary</i>	List 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 <i>-o</i> option to print the output to a file rather than the console.
<i>-username</i>	Specify the user name to log on to InterChange Server Express. If no user name is specified, <i>repos_copy</i> prompts for a user name.
<i>-v</i>	Print the version number of the program that the <i>repos_copy</i> utility executes.
<i>-vp</i>	Validate a package file. The server validates packages against the repository and makes sure that the dependencies among the components in the package are resolved. If the validation is not successful, <i>repos_copy</i> prints a list of the missing dependencies. This option does not make any changes to the repository; it just validates the package file. When using the <i>-vp</i> option you must also use the <i>-i</i> option to specify the package file to be validated.
<i>-vr</i>	Validate the repository. The output message indicates whether the validation is successful or not. If the validation is not successful, <i>repos_copy</i> prints a list of the missing dependencies.

Table 14. *Repos_copy* command options (continued)

Option	Description
-wi	Do not display warning messages 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 example.
-xCompileAll	Compile all collaboration templates and maps in the repository.
-xCompileAllCollabs	Compile all collaboration templates in the repository.
-xCompileAllMaps	Compile all maps in the repository.
-xCompileCollab: <i>collabTemplateName</i> [<i>+collabTemplateName</i>]	Compile the specified collaboration templates in the repository.
-xCompileMap: <i>nativeMapName</i> [<i>+nativeMapName</i>]	Compile the specified maps in the repository.
-xCompilePackage	Automatically compile the package deployed to the server. Because the production-mode server automatically compiles all packages, this option applies only to design-mode servers. For a full description of InterChange Server Express modes, see the System Implementation Guide.

Table 15. 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

Repos_copy usage scenarios

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

- “Example of printing the *repos_copy* command” on page 92
- “Example of validating a package” on page 92
- “Example of deploying a package to the repository” on page 92
- “Example of validating the repository” on page 94
- “Example of compiling components in the repository” on page 94
- “Example of deleting components from the repository” on page 94
- “Example of exporting components to a package” on page 96
- “Example of printing a list of components in the repository” on page 96

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 fails. 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.

For example:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1
-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 file named `WebSphereICS420DEVServer.jar` being deployed to a repository:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1
-iWebSphereICS420DEVServer.jar
```

Example of duplicate components: Commonly there are 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
```

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 Express 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 example, 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, you want to change the relevant properties before schemas are created. 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 example, `-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 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 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
```

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 and 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 fails. 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 when 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 15 on page 91. The following example deletes the relationship named **Customer**:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin  
-pnull -doRelationship:Customer
```

Examples 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.

Example of forcing a delete in spite of references: If a component has referents, `repos_copy` does 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 leaves 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
```

Example of 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 15 on page 91, 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 click **Show Dependencies**, 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 `-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 InterChange Server Express version levels earlier than 4.1.1 can be loaded into the repository correctly only if the dates and times for the component schedules are in full US format. (This is usually not a problem. Repos_copy saves all schedule dates in full US format only. The incompatibility may arise if the repos_copy files have been manually edited.)

In addition, when deploying 4.1.1 repositories, be sure to specify the character encoding with the '-nc' option.

Chapter 3. Troubleshooting the system

This chapter provides troubleshooting topics to help determine and resolve problem scenarios that may occur in an InterChange Server Express system. The following topics are covered:

“Using log and trace files for troubleshooting” on page 99

“Steps for controlling server memory usage” on page 121

“Maintaining event archives” on page 122

“Managing Work-in Progress connections” on page 123

“Managing failures” on page 123

“Working with failed events” on page 129

“Resolving collaboration deadlocks” on page 143

“Managing IBM WebSphere MQ” on page 147

“Steps for managing database connection pools” on page 148

“Steps for managing run-time properties” on page 148

“Managing the Object Request Broker” on page 149

Using log and trace files for troubleshooting

This section describes how to use log and trace files for troubleshooting. The following topics are covered:

“Overview of log and trace files” on page 99

“Configuring e-mail notification of log messages” on page 107

“Configuring logging and tracing” on page 110

“Working with tracing in Log Viewer” on page 115

Overview of log and trace files

Logging is used to communicate system messages, component state changes, failures, and tracing information. Messages that are generated from InterChange Server Express, collaboration objects, and connectors are sent to the destination you specified when you installed InterChange Server Express, by default, STDOUT (standard output). Messages that are generated from the connector agents are sent to STDOUT, but can be configured to be sent to a separate log file at the agent’s location. The messaging system is always active and provides an accurate monitor of the system.

Note: If you configure the messages generated from the connector agent to be sent to a separate log file, you must specify a log file or location that is separate from the InterChange Server Express log file.

You can configure the messaging system to send messages to a log file or an e-mail recipient in addition to the standard output. You can configure backup files (archives) for the log file, and determine their size. *Tracing*, which is disabled by default because of its effect on system resources, can be configured when problems arise and detailed information is needed for troubleshooting.

Two tools provide a graphical user interface for configuring and viewing message logging and tracing:

- The Edit Configuration tool opened from System Manager is used to set up or change system messaging and tracing for InterChange Server Express, collaborations, and business objects.
- LogViewer is used to display message and trace logs and to search the system message file for an explanation of a message.

In addition to using LogViewer to view logs, you can open the log with a text editor or create your own tools to filter the log file.

For information about viewing logging and tracing messages using Log Viewer, see “Overview of Log Viewer and tracing” on page 115.

This section covers the following topics:

“Collaboration object messages” on page 100

“Connector messages” on page 100

“Message formats” on page 101

“System logs” on page 102

“Log/trace file management” on page 103

Collaboration object messages

A collaboration object can generate messages to report run-time information, warnings, and errors. For example, a collaboration might log its decision points and the results of operations.

As InterChange Server Express executes collaboration objects, it writes their messages to its log. For information about configuring system logging, see “Configuring logging and tracing” on page 110.

In addition, you can send a collaboration object’s messages by e-mail to one or more recipients. You can specify a separate set of e-mail recipients for each collaboration object. For information on the rules for using e-mail notification, refer to “Configuring e-mail notification of log messages” on page 107.

Connector messages

Connector messages are sent to the InterChange Server Express message destination. A connector logs message to STDOUT by default, but can be configured to send a local destination log file to the InterChange Server Express logging destination.

For information about connector logging, see “Connector Agent logs” on page 102.

To aid in troubleshooting, a temporary log file is created during the connector agent boot-up that contains metadata obtained from the connector controller. The metadata consists of business object specifications, properties, and delta-supported properties. The file is named *connectornametmp.log* and is found in the *ProductDir\Connectors* directory.

Message formats

All messages are formatted so they can easily be filtered. Logged messages for InterChange Server Express and connectors use the same format, which is described in Table 16.. When business objects are configured for flow tracing messages, they use these fields and the additional fields, denoted by an asterisk in Table 16.. A message delivered to InterChange Server Express has the following format, using some or all of the following parameters:

Time: System Name: Thread: MsgType MsgID: SubSystem: FIID: BO: MsgText: BOD:

Table 16. Message format

Variable	Description
<i>Time</i>	Timestamp: the date of logging in the format <i>year/month/date time</i> .
<i>System</i>	Type of component (system identifier). It can be Server, Collaboration, Business Object, or ConnectorAgent.
<i>Thread</i>	Thread name and thread ID
<i>Name</i>	Name of the component, such as ClarifyConnector.
<i>MsgType</i>	Severity of the message. See Table 17.
<i>MsgID</i>	Message number.
<i>SubSystem*</i>	Subsystem of the current system. It can be Event Management, Messaging, Repository, or Database Connectivity.
<i>FIID*</i>	Flow initiator ID of the business object.
<i>BO*</i>	Business object name.
<i>MsgText</i>	Associated text for the message number.
<i>BOD*</i>	Business object dump. The data contained in the business object.

Following is an example of a message for the server: [Time: 2001/06/07 11:01:29.487] [System: Server] [SS: REPOSITORY] [Thread: VBJ ThreadPool Worker (#-1767149274)] [Type: Trace] [Mesg: Released session REPOSITORY0]

Table 17 describes the types of InterChange Server Express messages.

Table 17. Message types

Type	Description
Info	Informational only. You do not need to take action.
Warning	A default condition chosen by InterChange Server Express.
Error	A serious problem that you should investigate.
Fatal Error	An error that stops operation and should be reported.
Trace	Tracing information for the trace level specified.
Flow Trace	Flow tracing information for business objects.
Internal Error	A serious internal problem that should be investigated.
Internal Fatal Error	An internal error that stops operation. It should be reported.

Note: If a message with the Internal Error or Internal Fatal Error severity appears, record the circumstances surrounding the problem, and then contact InterChange Server Express Technical Support.

System logs

InterChange Server Express logs: InterChange Server Express can log messages to the following destinations:

- The default system log file, `InterchangeSystem.log`, which you can create in the product directory.
- A log file that you specify.
- Standard output (STDOUT). STDOUT messages appear in the Command Prompt window in which InterChange Server Express starts.
- In addition to managing regular logging, InterChange Server Express can send an e-mail message to a specified user when it generates error or fatal error messages.

By default, trace messages are sent to the system log file. In some cases, this file may become too large, so it is recommended that you specify a separate trace file. See “Trace files” on page 104 for instructions on setting up a separate file for trace messages.

“Configuring logging and tracing” on page 110 describes how to set the destination for logging.

To specify the recipient for e-mail notification, see “Configuring e-mail notification of log messages” on page 107.

“Log/trace file management” on page 103 describes how to keep log files from becoming too large.

Connector Agent logs: The connector agent and connector controller have separate mechanisms for logging. This section describes connector agent logging. Connector controller messages are sent to the log that contains the InterChange Server Express messages.

A connector agent logs messages to a local destination and can also send its messages to InterChange Server Express for logging. To specify a log file name, edit the `LogFile` property of the connector (using the Connector Configurator) and insert the name of the log file you want to use.

If you want to specify a logging and tracing file for the connector agent, you must specify a local configuration file when starting the agent.

The default log file (located at STDOUT), contains text for the error and informational messages raised from the connector. *Name* is the name of the application.

Table 18 on page 103 describes the properties you can edit that determine where a connector agent logs messages.

Table 18. Connector agent log message properties

Property name	Description	Type of value
LogAtInterchangeEnd	Specifies whether the connector agent sends messages to InterChange Server Express in addition to logging them locally. At InterChange Server Express, connector agent messages appear wherever server messages appear, according to the InterchangeSystem.cfg file.	Either true (sends messages to InterChange Server Express and enables e-mailing) or false (logs messages only locally). The default value is false.
LogFileNames	Specifies where to write connector agent messages on the local system.	A file path or standard output (STDOUT). The default value is STDOUT.

For instructions on configuring these properties, see “Steps for configuring the connector agent logging destination” on page 112.

Log/trace file management

When the InterChange system is started, a log file is created if one does not exist, or is appended if it does. If the size of the log file is unlimited, it grows and its size depends on the amount of time since it was last managed and the volume of transactions passing through the system. If a log file grows too large, you may not be able to open it or an application may require additional system resources to write to the files.

InterChange Server Express system log files can be configured to a specified size and then automatically archived after they reach that size. As an added precaution, you can specify a number of archive files to use as a system backup. Each time the log file reaches its maximum size, the file is renamed as a new archive file. The archive file’s name is derived from the original log or trace file name, with the following inserted into the name:

Arc number

For example, using 5 archive files, if the log file has the name InterchangeSystem.log, the first archive created is named InterchangeSystem_Arc_01.log. When the new log file fills up, InterchangeSystem_Arc_01.log is renamed InterchangeSystem_Arc_02.log, and the log file is again saved to InterchangeSystem_Arc_01.log and so on in a circular fashion, until there are five archive files. If there are five archive files, when a new log file is created, existing archive files are renamed and their numbers incremented so the number of archives matches the number you configured, then the oldest file, whose archive number is 05, is deleted. Figure 48 on page 104 shows the progression of files using this configuration.

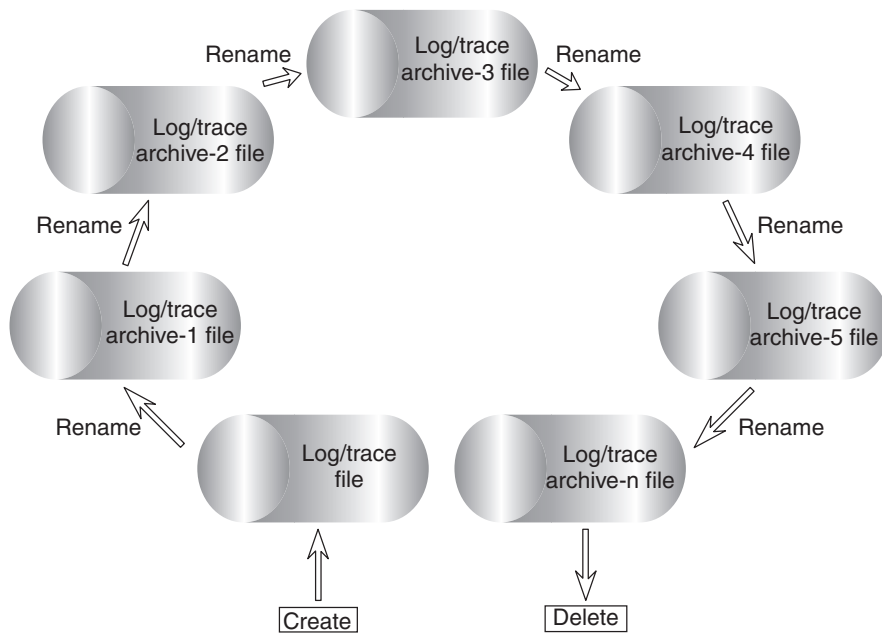


Figure 48. Circular archival logging

See the configuration tasks “Configuring logging and tracing” on page 110 for details.

If the system log file is configured for unlimited size, InterChange Server Express writes to the log until the disk that the log file is located on gradually fills; if the disk is not managed, the system returns an error message when the disk is full.

The data in the following files should be deleted periodically:

- InterChange System log file. Contains a record of all interactions of the various system components with InterChange Server Express. The Edit Configuration tool allows you to specify the name and size of this file and a number of archive files for automatically saving versions of this file when it overruns its specified size. See “Steps for configuring InterChange Server Express logging and tracing destinations” on page 110 for task information.
- Connector log file. Contains flows at the API level that were executed by the connector. Not all applications support these files.
- Connector error log file. Contains error messages returned by the connector. Not all applications support these files.

In addition to the log files above, other log files exist that are specific to each application. Most files are created during run time if they do not already exist. New information is appended to any existing file. Each component that supplies log information to the files must be stopped before proceeding with a backup.

Any file management procedure can be used, but the following periodic log file management is recommended:

- Rename the files by appending a date to the file.
- Move the files to an archive directory.

Trace files

To troubleshoot a problem, you can turn on tracing. *Trace messages* help you monitor actions taken in components of the InterChange Server Express system.

Trace levels define the amount of detail written to the trace file. The higher the trace level, the more detail you receive. Tracing differs from logging in the following ways:

- Logging always occurs, but tracing can be turned on and off as needed.
- Tracing contains more detailed information than logging about the state of components and the actions taken by them.
- Logging and tracing settings are persistent after reboots.

Tracing is off by default because it produces messages that are more detailed than you normally need. You can turn tracing on and off as necessary when InterChange Server Express is running.

Tracing services of InterChange Server Express: Tracing services for InterChange Server Express are initially set in parameters of the configuration file for InterChange Server Express (by default, this file is called `InterchangeSystem.cfg` and resides in the product's top level directory. For details about these parameters, refer to *WebSphere Business Integration Server Express Installation Guide*. Settings for these parameters can be updated in the Edit Configuration tool of System Manager, as described in "Steps for configuring tracing levels for InterChange Server Express, business objects, and collaborations" on page 113.

Collaboration object tracing: You can trace the execution of a collaboration object. Tracing writes detailed messages about execution of the collaboration object to the log destination, which is specified in the `InterchangeSystem.cfg` file. Tracing collaborations is persistent. There are two trace level settings for collaborations, system level and collaboration level.

System level tracing returns run-time information for the collaboration. For example, if you want to trace the state changes of the collaboration, set the system trace level to 3.

You can set collaboration object tracing to one of the following levels:

Table 19. Collaboration object tracing levels

Trace level	Description
1	Traces the receipt of business objects from connectors and the start of the appropriate scenarios.
2	Prints messages for level 1. In addition, traces the start and completion of each scenario, reporting both forward execution and rollback.
3	Prints messages for levels 1 and 2. In addition, traces the execution of each scenario decision block or action.
4	Prints messages for levels 1 through 3. In addition, traces the sending and receipt of each business object by each scenario.
5	Prints messages for levels 1 through 4. In addition, traces the sending and receipt of each business object by each scenario, printing the value of each attribute in the business object.

For configuration instructions, see "Steps for configuring tracing levels for InterChange Server Express, business objects, and collaborations" on page 113.

Connector tracing: A connector contains two components, the connector controller and the connector agent. The two components can be in different locations on the network and are traced differently.

You can set connector agent and controller tracing to one of the following levels:

Table 20. Connector agent and controller tracing levels

Trace level	Description
1	Traces initialization and the sending and receipt of business objects.
2	Prints messages for level 1. In addition, provides more details than Level 1 for the same types of events.
3	Prints messages for levels 1 and 2. In addition, traces the exchange of messages between the connector agent and the messaging driver.
4	Prints messages for levels 1 through 3. In addition, traces the passing of business objects between internal levels of the connector.
5	Prints messages for levels 1 through 4. In addition, traces the passing of administrative messages between internal levels of the connector.

A new or changed tracing level takes effect immediately.

If you want to specify a logging and tracing file for the connector agent, you must specify a local configuration file when starting the agent.

For configuration instructions, see “Steps for configuring connector tracing” on page 114.

Map tracing: You can use System Manager to set up tracing for maps. Tracing maps is useful for debugging and keeping track of information and error messages created by the map. Tracing of maps is turned off by default.

For more information about tracing maps, see the *Map Development Guide*. For configuration instructions, see “Steps for configuring map tracing” on page 115.

Business object tracing (flow tracing): Business object trace logging provides a way to trace the progression of business objects from one processing point to another, based on notification messages that are generated at each point. For example, with level 2 tracing, when a business object arrives at a collaboration for processing, a trace message is logged.

Table 21 describes the configurable levels associated with business object tracing:

Table 21. Business object tracing levels

Trace Level	Description
0	No tracing.
1	Event status (such as Successful or Failed) and event identity information.
2	Minimal event tracing. Information about when a business object enters/exits systems, such as connectors, maps, relationships, and collaborations. Includes level 1 information.
3	Provides event tracing from level 2 and a business object dump at entry/exit of systems. System performance impact.
4	Detailed tracing. Provides tracing for system components such as connectors, maps, relationships, and collaborations, and maps these traces to level 3 event tracing. This level of tracing affects system performance.

For configuration instructions, see “Steps for configuring tracing levels for InterChange Server Express, business objects, and collaborations” on page 113.

Web gateway tracing: Web gateway tracing is provided at two levels, minimum and maximum. This tracing allows you to view information pertaining to whether

communication processing is being performed correctly on the gateway. The gateway trace level is set from the Gateway Configuration Tool. For instructions, see the *System Implementation Guide*.

Table 22 describes the configurable levels associated with gateway tracing:

Table 22. Gateway tracing levels

Trace Level	Description
0	No tracing.
1	Minimum tracing (such as bind requests, socket openings, and so forth).
5	Maximum tracing (such as HTTP/HTTPS requests, including all headers).

Configuring e-mail notification of log messages

Error and fatal messages that are logged to the InterChange Server Express log can also be sent to the InterChange Server Express system administrator, or any other recipient, by e-mail. By default, InterChange Server Express is configured to send e-mail notifications using JavaMail, but you can configure the server to send e-mail notifications using the e-Mail connector. For instructions on configuring e-mail notification using the e-Mail connector, refer to “Steps for configuring e-mail notification at the system level” on page 108.

Note: If you want to configure collaborations for e-mail notification, you must use the e-Mail connector.

The following components can be configured for sending error and fatal messages to an e-mail recipient:

- InterChange Server Express
- Collaboration objects
- Connectors

You can configure e-mail notification at the system level (set in the `InterchangeSystem.cfg` file), at the collaboration object level (set as a collaboration object property), or at the connector level (set as a connector property). If you configure e-mail notification at the system level, the configuration applies to all of the collaboration objects or connectors in the system. If you configure e-mail notification at the collaboration object or connector level, the configuration applies only to that specific component and supersedes the system configuration.

E-mail notification in the InterChange Server Express system supports Simple Mail Transport Protocol (SMTP) mail messages, therefore, the e-mail recipient value in the `InterchangeSystem.cfg` file and the collaboration e-mail addresses must be standard Internet addresses.

A valid e-mail address entry can be one or more fully qualified Internet addresses, separated by commas. For example, a valid entry for two recipients is:

```
JohnDoe@company.com,FredSmith@company.com
```

You cannot use personal address aliases, such as an alias defined in a personal address book. However, a valid address can be an alias defined in a mail server, such as `Eng@company.com`. In this case, the mail server decodes the alias and sends e-mail to all members of the alias. For example, a decoded alias might be `person1@some_company.com`, `person2@another_company.com`, and so forth.

Steps for configuring e-mail notification at the system level

Perform the following steps to configure e-mail notification at the system level:

1. Open System Manager
2. In the InterChange Server Component Management view, right-click the server and click **Edit Configuration**.

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

3. Click the **E-mail** tab.

A dialog box appears in the upper-right section of the System Manager window in which you can enter the parameters necessary for configuring e-mail notification at the system level (see Figure 49).

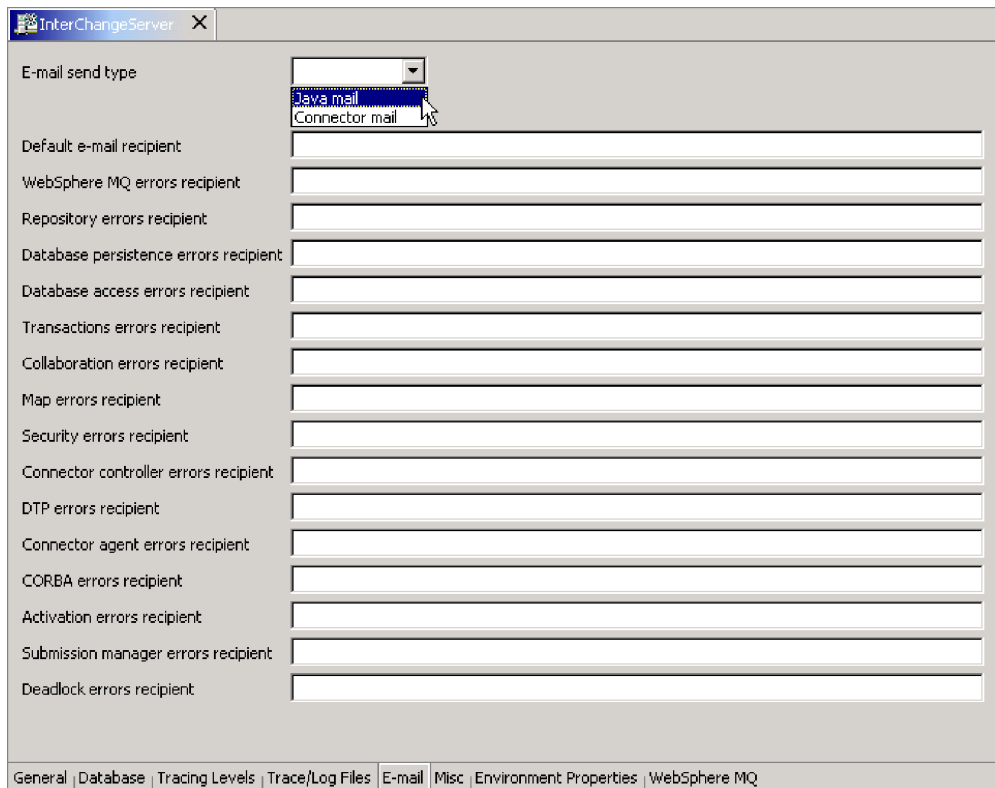


Figure 49. Edit Configuration, E-mail tab

4. Select an e-mail type from the **E-mail send type** list:
 - Select **Java mail** if you do not plan to configure collaborations for e-mail notification. An **SMTP mail host** field appears at the top of the list of subsystems. Type the host name of the SMTP mail server.
Note: Be sure to use the correct case, as this field is case sensitive.
 - Select **Connector mail** if you plan to configure collaborations for e-mail notification. If you use this type of e-mail you must install the e-Mail connector, then configure the e-Mail connector and e-Mail business object. For instructions on installing the e-Mail connector, see *WebSphere Business Integration Server Express Installation Guide*. For instructions on configuring the e-Mail connector and e-Mail business object, see steps 7 and 8.
5. Type one or more valid e-mail addresses in each subsystem field. The address must be SMTP-compliant.

6. In System Manager, click **File > Save ServerName**. The system-level e-mail notification information you entered is saved in the InterchangeSystem.cfg file.
7. Additional step if you configured **Connector mail**: Configure the EmailNotification business object to hold the e-mail address of the person who receives e-mail if the intended e-mail recipient is unreachable. As a fail-safe, this should probably be the mail administrator, not the InterChange Server Express system administrator, to ensure the mail is delivered if the InterChange Server Express system administrator is unreachable.

Do the following:

- In System Manager, right-click the EmailNotification business object, and click **Edit definition**. The EmailNotification Business Object Designer Express window appears.
 - On the **Attributes** tab, type the return e-mail address in the **Default** column of the FromAddress attribute.
 - Click **File > Save**.
 - Close Business Object Designer Express.
8. Additional step if you configured **Connector mail**: Configure the e-Mail connector.

Steps for configuring e-mail notification at the collaboration object level

Perform the following steps to set an e-mail address to receive messages for a collaboration:

1. In System Manager, go to the InterChange Server Component Management view, right-click the collaboration object for which you want to configure e-mail notification, and click **Properties**.
The Properties dialog box appears.
2. In the **Collaboration General Properties** tab, type a valid e-mail address in the **Email notification address** field. The address must be SMTP-compliant.
3. Click **OK** to save your changes and close the window.
4. Restart the collaboration for changes to take effect.

Steps for configuring e-mail notification at the connector level

You can configure e-mail notification for a specific connector, however, configuration parameters set at the collaboration object level supersede those set at the system level.

Perform the following steps to configure e-mail notification for a specific connector:

1. In System Manager, right-click the connector for which you want to configure e-mail notification, and click **Edit Definition**.
The Connector Configurator Express window appears.

2. In the **Standard Properties** tab, select true in the **Value** field of the LogAtInterchangeEnd property. This enables connector messages to be mailed to the InterChange Server Express log.
3. Click **File > Save > To Project**.
4. Close the Connector Configurator Express window.
5. Restart the connector for the change to take effect.

Configuring logging and tracing

This section describes how to set up message logging and tracing. These settings can be made by using The Edit Configuration option in System Manager and by manually editing the InterchangeSystem.cfg file. The following tasks are described:

“Steps for opening the Edit Configuration tool” on page 110

“Steps for configuring InterChange Server Express logging and tracing destinations” on page 110

“Steps for configuring other InterChange Server Express logging and tracing parameters” on page 112

“Steps for configuring the connector agent logging destination” on page 112

“Steps for configuring tracing levels for InterChange Server Express, business objects, and collaborations” on page 113

“Steps for configuring the collaboration object trace level” on page 114

“Steps for configuring connector tracing” on page 114

“Steps for configuring map tracing” on page 115

Steps for opening the Edit Configuration tool

InterChange Server Express must be running and in the Connected state to use the Edit Configuration tool. The Edit Configuration tool can manage only one InterChange Server Express per session.

Perform the following steps to open the Edit Configuration tool:

1. Open System Manager.
2. Right-click a server instance in the InterChange Server Component Management view and click **Edit Configuration**. The Edit Configuration tool opens in the upper-right section of System Manager, where you can edit the InterchangeServer.cfg file in the

Steps for configuring InterChange Server Express logging and tracing destinations

Perform the following steps to configure the destination for InterChange Server Express message logging and tracing:

1. In the Edit Configuration tool, click the **Trace/Log Files** tab.

The logging and tracing configuration window appears (see Figure 50 on page 111).

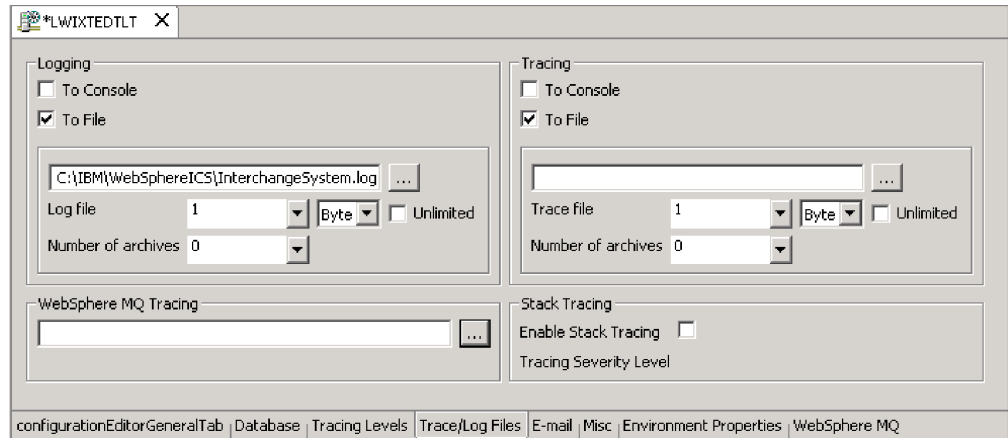


Figure 50. Edit Configuration tool, Trace/Log Files tab

2. Under **Logging**, select the destination for system logging. You can log to both system console and a log file, but use this option only for debugging and only on development systems. Do the following:
 - To log to the system console (standard output), which is the default, select the **To console** (STDOUT) check box. To disable the console as the logging destination, clear the **To console** (STDOUT) check box.

Note: If you plan to run InterChange Server Express as a service, you must set logging and tracing to file destinations. Setting logging and tracing to STDOUT prevents InterChange Server from being configured as a Windows service.

 - To log to a file, select the **To File** check box and either type in the full path name of the file or click the browse button to navigate to the log file.
 - If you select a log file, you can configure the size of the file in MB (or keep the default, **Unlimited**) and the number of archive files to create. For information about archives, see “Maintaining event archives” on page 122.
3. Under **Tracing**, select the destination for system tracing. Do the following:
 - To log to the system console (standard output), select the **To console** (STDOUT) check box. To disable the console as the logging destination, clear the check box.

If you want to log messages to STDOUT, the message appears in the Command Prompt window in which InterChange Server Express starts.

 - To log to a file, select the **To File** check box and type in the full path name of the file or click the brows button to navigate to the trace file.
 - If you select a trace file, configure the size of the file in MB (or keep the default, **Unlimited**, selected) and the number of archive files to create. For information about archives, see “Maintaining event archives” on page 122.
4. To configure IBM WebSphere MQ tracing, accept the default path name of the file or click the browse button to navigate to the log file.
5. Click **OK** to save changes and exit.

Your changes take effect immediately, and if you already had a log file configured, it is saved and dated, and a new file created using the newly entered configuration.

Steps for configuring stack trace

When InterChange Server Express stack tracing is enabled, the stack trace information is printed to the InterChange Server Express log file. This can be useful for troubleshooting a newly installed system.

Perform the following steps to configure the stack tracing feature:

1. Open the Edit Configuration tool (see “Steps for opening the Edit Configuration tool” on page 110).
2. Click the **Trace/Log Files** tab.
3. Select the **Enable Stack Tracing** check box. When you enable stack tracing, the **Tracing Severity Level** list is available.
4. In the **Tracing Severity Level** list, select the severity level for which you want to print the stack tracing. The choices are: **INFO**, **WARNING**, **ERROR**, and **FATAL**. When stack tracing is printed to the log file, it prints the level you select and all levels above it. For example, if you select **ERROR**, both **ERROR** and **FATAL** messages are printed to the log file.

Steps for configuring other InterChange Server Express logging and tracing parameters

Perform the following steps to set file size and number of archived files for log and trace files:

1. Open the Edit Configuration tool (see “Steps for opening the Edit Configuration tool” on page 110).
2. Click the **Trace/Log Files** tab.
3. To limit the log file, select the **To File** and the **Unlimited** check boxes, and select a number to represent the maximum log file size from the **Log file** list, and then select **Byte**, **KB** (kilobyte), **MB** (megabyte), or **GB** (gigabyte).
4. To limit the trace file, select the **To File** and the **Unlimited** check boxes, select a number to represent the maximum trace file size from the **Trace file** list, then select **Byte**, **KB** (kilobyte), **MB** (megabyte), or **GB** (gigabyte).
5. To set the number of log or trace file archives, select a number from the **Number of archives** list under either **Logging** or **Tracing** section.
6. Stop and restart InterChange Server Express.

Steps for configuring the connector agent logging destination

The two connector components have separate mechanisms for logging. Connector controller logging is sent to the `InterchangeServer.log` file. This section describes configuring the connector agent log file name and location.

For background information on connector agent logging, see “Connector Agent logs” on page 102.

You can set the destination for connector agent logging using one of the following methods:

- If you are installing Remote Agent, you can configure the logging destination during the installation. For more information about installing Remote Agent, see the *System Implementation Guide*.
- If you are not using Remote Agent, you can use Connector Configurator Express to configure the connector agent logging destination.

Perform the following steps to configure the destination for connector agent logging using Connector Configurator Express:

1. Connect to a server.

2. In System Manager, right-click a connector object, then select **Edit definition**. The Connector Configurator Express window appears.
3. Click the **Standard Properties** tab.
4. Select one of the logging property values (see Table 18 on page 103 for an explanation of these values), and click **Edit**.
5. Type the new value in **Value** field and click **OK**.

For example, change the LogAtInterchangeEnd value to true to send messages to the InterChange Server Express log. If InterChange Server Express is configured to send e-mail when error and fatal messages are logged, e-mail is sent for the connector agent messages as well.

6. Repeat Steps 3 and 4 to edit the other logging property value if necessary.
For example, to send connector agent messages to a message file instead of the default STDOUT, enter the full path name of the file in the **Value** field.
7. Restart the connector for changes to take effect.

Note: If you want to specify a logging and tracing file for the connector agent, you must specify a local configuration file when starting the agent.

Steps for configuring tracing levels for InterChange Server Express, business objects, and collaborations

Perform the following steps to configure tracing:

1. Open the Edit Configuration tool (see “Steps for opening the Edit Configuration tool” on page 110).
2. Click the **Tracing Levels** tab (see Figure 51).

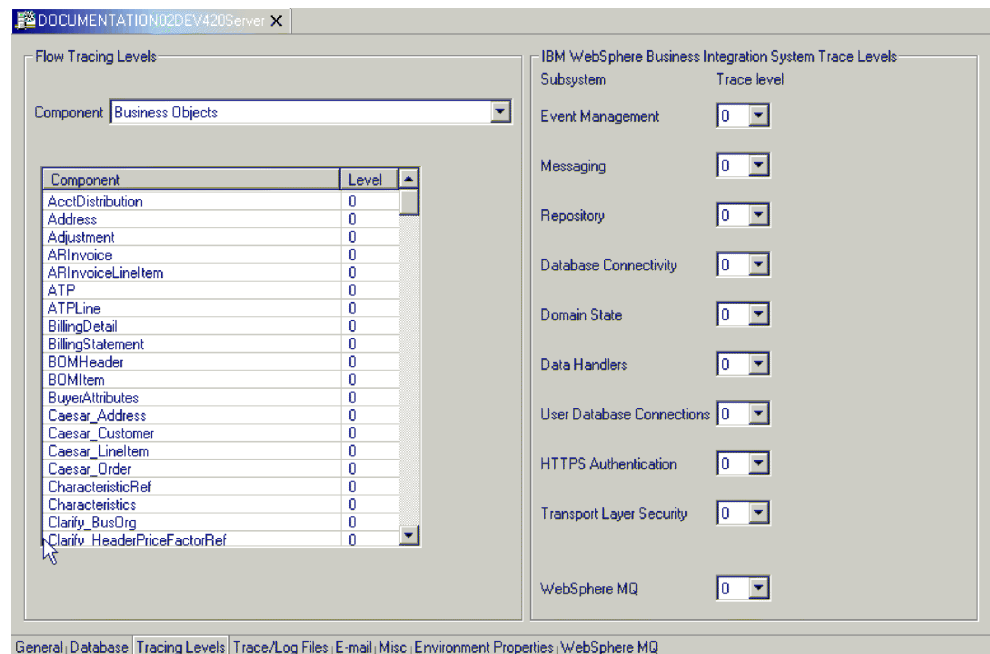


Figure 51. Edit Configuration tool, Tracing Levels tab

3. To configure subsystem tracing, under **IBM WebSphere Business Integration Trace Levels**, select a **Trace level** for the subsystems you want to configure. Setting these trace levels updates parameters in the InterchangeSystem.cfg file. For details about what information is produced at the various tracing levels, see *WebSphere Business Integration Server Express Installation Guide*.

4. To configure tracing for collaborations, do the following:
 - Under **Flow Tracing Levels**, select **Collaborations** from the **Component** list. The collaboration names that are configured for the system appear in the table.
 - For each component, select a level for tracing in the **Level** column. For a description of the trace levels for collaborations, see “Collaboration object tracing” on page 105.
5. To configure tracing for business objects, do the following:
 - Under **Flow Tracing Levels**, select **Business Objects** from the **Component** list. The business object names that are configured for the system appear in the table.
 - For each component, select a level for tracing in the **Level** column.
6. Click **OK** to save changes and exit.

The trace levels for the subsystems, business objects and collaborations are immediately in effect.

Steps for configuring the collaboration object trace level

This section describes one of two methods that can be used to configure collaboration object trace levels. For the alternative method, see “Steps for configuring tracing levels for InterChange Server Express, business objects, and collaborations” on page 113.

Perform the following steps to configure the collaboration object trace level:

1. In System Manager, go to the InterChange Server Component Management view, right-click the collaboration object, and click **Properties**. The Properties dialog box appears.
2. Do one of the following on the Collaboration General Properties tab:
 - If you want to configure run-time system trace level for a collaboration object, select a **System trace level** value, then click **OK**. Collaboration object tracing returns messages from inside the collaboration. For a description of collaboration object trace levels, see “Collaboration object tracing” on page 105.
 - If you want to configure collaboration object tracing, select a value in the **Collaboration Trace Level** list, and click **OK**.

A collaboration object starts tracing as soon as its tracing level changes.

Steps for configuring connector tracing

Perform the following steps to configure connector tracing:

1. From System Manager, right-click the connector object, then click **Edit definition**. The Connector Configurator Express window appears.
2. Do one of the following:
 - To set a connector controller’s tracing level, from the Standard Properties tab, click in the **Value** field of the **ControllerTraceLevel** property, then set the controller tracing level.
 - To set a connector agent’s tracing level, from the Standard Properties tab, click the **Value** field of the **AgentTraceLevel** property, then set the agent tracing level.
3. Click **File > Save > To Project**.
4. Close Connector Configurator Express.

Any changes to these properties take effect immediately. Trace messages for connector controllers appear wherever InterChange Server Express sends its trace messages, and trace messages for connector agents appear wherever the connector agent logs messages.

Note: If you want to specify a logging and tracing file for the connector agent, you must specify a local configuration file when starting the agent.

Steps for configuring map tracing

Perform the following steps to set the trace level for a map:

1. In System Manager, go to the InterChange Server Component Management view, right-click the map object, and click **Properties**. The Maps Property Page appears (see Figure 52).

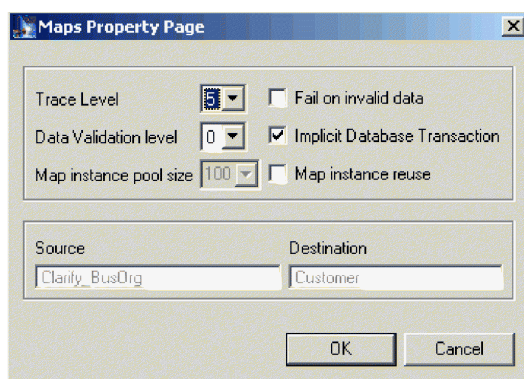


Figure 52. Maps Property Page

2. In the **Trace Level** field, type the appropriate value.
3. Click **OK**.

Working with tracing in Log Viewer

To troubleshoot a problem, you can turn on tracing. Trace messages help you monitor actions taken in components of the WebSphere InterChange Server Express system. Trace levels define the amount of detail written to the trace file. The higher the trace level, the more detail you receive.

Tracing is off by default because it produces messages that are more detailed than you normally need. You can turn tracing on and off as necessary when InterChange Server Express is running.

For background information about tracing, see “Trace files” on page 104. For configuration information about tracing, see “Configuring logging and tracing” on page 110.

Overview of Log Viewer and tracing

Use LogViewer to display trace information for InterChange Server Express. In addition to displaying the debugging trace information for collaborations, it allows you to view the progression of a business object as it passes from one processing point to another, for example as the business object exits the collaboration processing point and is sent on to other collaborations or connectors for processing or is forwarded to the mapping stage for data transformation. For information about flow tracing, see “Business object tracing (flow tracing)” on page 106.

InterChange Server Express system message logging is used to communicate messages, component state changes, and failures.

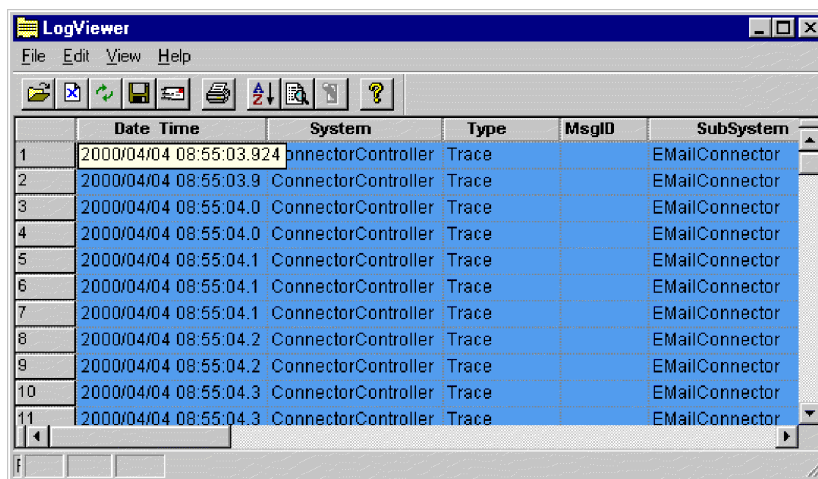
Note: Log Viewer views log files and therefore does not need InterChange Server Express to be running. The InterChange Server Express system administrator must have the appropriate file system permissions set to view log files.

You can view log files containing messages and explanations of system messages either of these ways:

- Use Log Viewer to view system log files in its window, and obtain a message explanation interactively by clicking on a message number. See “Overview of Log Viewer and tracing” on page 115.
- Use a text editor to view a log file and to search the InterChange Server Express message file by message number or text content to obtain an explanation.

Log Viewer allows you to see all messages contained in a log file. You can sort and filter the output display and print, save, and e-mail of the file.

A sample log file is shown displayed in Log Viewer in Figure 53..



The screenshot shows the Log Viewer application window with a menu bar (File, Edit, View, Help) and a toolbar. The main area contains a table with the following data:

	Date Time	System	Type	MsgID	SubSystem
1	2000/04/04 08:55:03.924	ConnectorController	Trace		EEmailConnector
2	2000/04/04 08:55:03.9	ConnectorController	Trace		EEmailConnector
3	2000/04/04 08:55:04.0	ConnectorController	Trace		EEmailConnector
4	2000/04/04 08:55:04.0	ConnectorController	Trace		EEmailConnector
5	2000/04/04 08:55:04.1	ConnectorController	Trace		EEmailConnector
6	2000/04/04 08:55:04.1	ConnectorController	Trace		EEmailConnector
7	2000/04/04 08:55:04.1	ConnectorController	Trace		EEmailConnector
8	2000/04/04 08:55:04.2	ConnectorController	Trace		EEmailConnector
9	2000/04/04 08:55:04.2	ConnectorController	Trace		EEmailConnector
10	2000/04/04 08:55:04.3	ConnectorController	Trace		EEmailConnector
11	2000/04/04 08:55:04.3	ConnectorController	Trace		EEmailConnector

Figure 53. Log Viewer

Steps for starting Log Viewer

Perform one of these steps to start Log Viewer:

- Select **Start > Programs > IBM WebSphere Business Integration Express > Toolset Express > Administrative > Log Viewer**. Click **File > Open** to browse for the log file.
- Use the Run command from the Start menu and browse for the LogViewer.exe file. Click **File > Open** to browse for the log file.

Steps for setting Log Viewer preferences

Perform the following steps to set Log Viewer preferences:

1. Click **Edit > Preferences**.

The User configuration options dialog box appears, with the **General** tab selected (see Figure 54 on page 117).

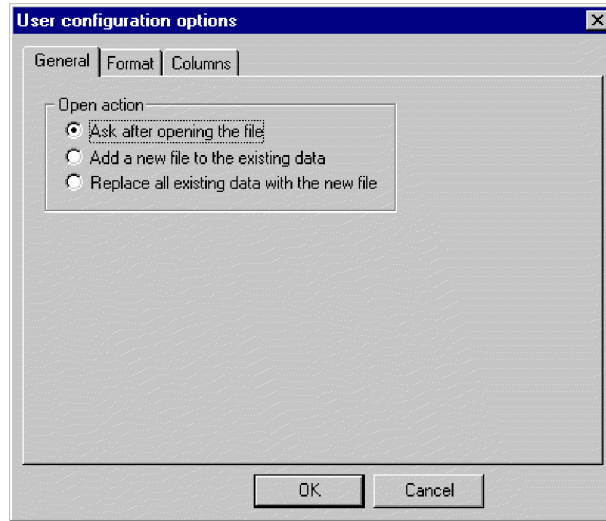


Figure 54. Log Viewer User configuration options dialog box, General Properties tab

2. Under **Open action**, select how to display the log file when you open a log file:
 - **Ask after opening the file:** You are asked what you want to do each time you open a log file.
 - **Add a new file to the existing data:** The opened log file is merged with the currently displayed log file.
 - **Replace all existing data with the new file:** The currently displayed log file is replaced with the contents of the opened log file.
3. To change the background color and font of the Log Viewer messages, click the **Format** tab.

The Format tab appears (see Figure 55).

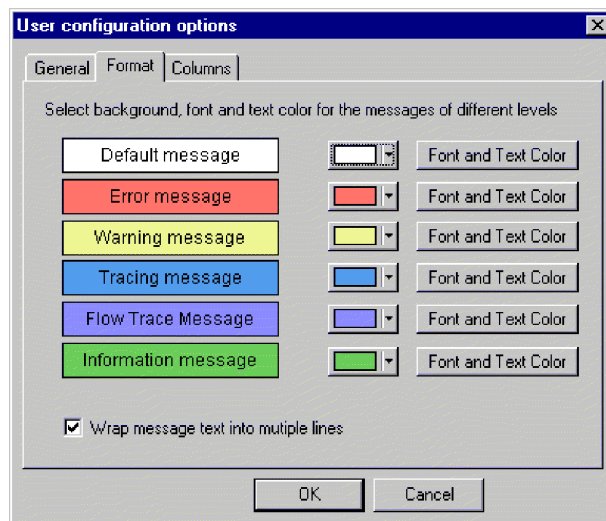


Figure 55. Log Viewer User configuration options, Format tab

4. Select options to display the various types of log messages:
 - Assign different background colors and fonts for each type of message that appears so you can easily recognize its severity (for example, a red background with larger font allows for Warning messages).

- Select the **Wrap message text into multiple lines** check box if you want the text to wrap when wider than the column.
5. To change which Log Viewer columns are displayed, click the **Columns** tab. The Columns tab opens (see Figure 56).

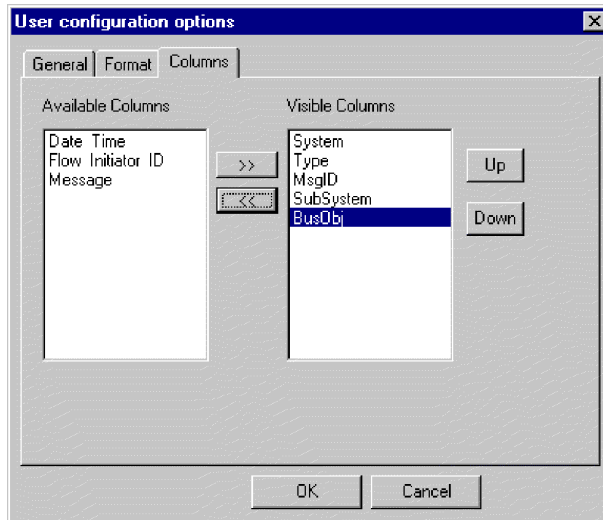


Figure 56. Log Viewer User configuration options, Columns tab

This window allows you to determine which columns display in Log Viewer:

- To display a column, select a column name in the **Available Columns** list and click the **>>** button to move it to the **Visible Columns** list.
- To hide a column, select a column name in the **Visible Columns** list and click the **<<** button to move it to the **Available Columns** list.
- Click any of the column names in the **Available Columns** list and click the **Up** or **Down** button to change its ordering from left to right in the Log Viewer. **Up** moves columns to the left and **Down** moves columns to the right.

Overview of Log Viewer display options

The **View** menu contains additional options to change how Log Viewer appears:

- Display/hide the Log Viewer toolbar.
- Display/hide the Log Viewer status bar.
- Split the window into two or more views.
- Sort the messages. Figure 57 on page 119 shows the Sort dialog box.

Click the down arrow in each sort field to select **Date/Time**, an EventID, or a business object. Further sort by ascending or descending order.

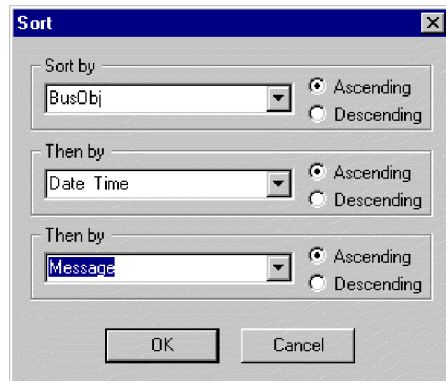


Figure 57. Log Viewer Sort dialog box

- Filter or show all messages. Configure filter options by checking filtering options in the filter tabs, such as time range or type of message (see Figure 58 and Table 17 on page 101). Under **Activate Filters**, select the check box that is associated with the tab where you selected the filter options, and click **OK** to enable filtering. The filtered output can be toggled on or off with the **Filter Toggle** button on the Log Viewer toolbar.

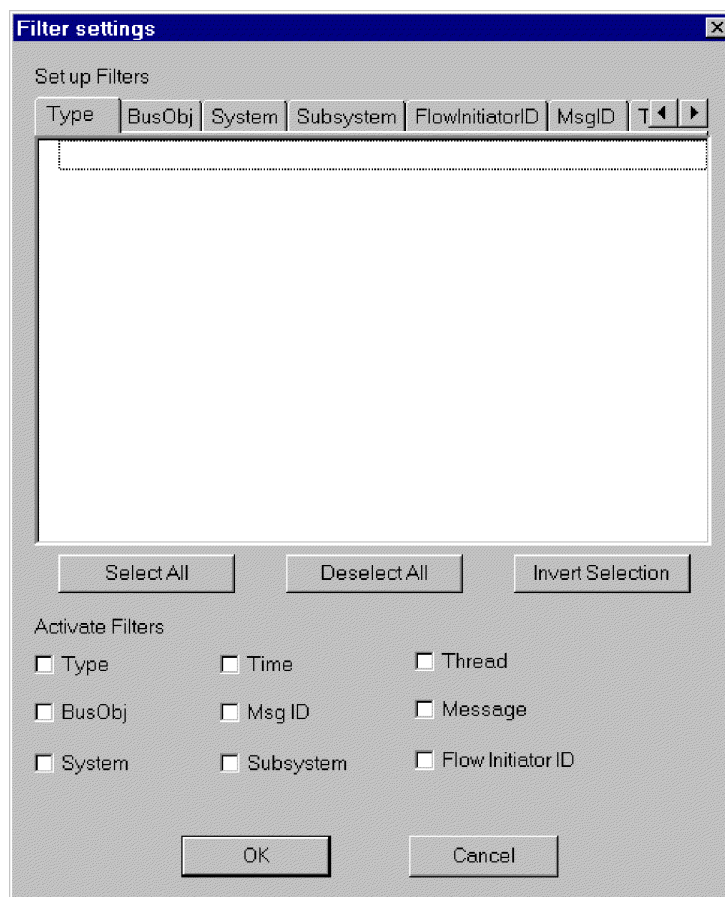


Figure 58. Log Viewer Filter Settings dialog box

Overview of Log Viewer display output

Several options are available for manipulating Log Viewer output. In the **File** menu, there are options for print previewing, printing, saving, refreshing the

display, sending to an e-mail recipient, and determining the style for page setup, headers, and footers. The variables for headers and footers are:

\$F	Name of file
\$A	Application name
\$P	Page number
\$N	Total number of pages
\$D	Date (can be followed by additional parameters (for example \$D{%y:%h:%m}))

Steps for filtering messages

Perform the following steps to filter messages in Log Viewer:

1. Click **View > Filter > Use Filter**. The Filter Settings dialog box appears. The Filter Settings dialog box displays categories that correspond to the parameters of the logging message format (see Table 17 on page 101 for a descriptive list of the parameters).
2. Select a tab under **Set up Filters** to display the items that you want to use for filtering messages. For example, select the BusObj tab to display a list of business objects to be used in filtering, and select Time if you want to filter according to the timestamp of the message. You can set up multiple filters, and use them either separately or with one another.
3. In the displayed list of items, select each item for which you want to view messages in Log Viewer. For example, if you want to view messages related only to the Cost and Customer business objects, select only those business objects in the list. If you want to view only messages that are time-stamped between 5 March 2002 at 9:00 a.m. and 6 March 2002 at 5:00 a.m., select the range for those times under the **Time** tab.
You can use the buttons below the list box to select all the displayed items, to clear all the displayed items, or to invert your current selected choices.
4. Under **Activate Filters**, select the check box for each filter type that you want to activate. For example, if you want to see all messages for the Cost and Customer business objects (which you specified in the previous step), activate only the BusObj filter. If you want to see only those messages for the Cost and Customer business objects that have a particular timestamp, activate both the BusObj filter and the Time filter.
5. Click **OK**. The Filter Settings dialog box closes, and the Log Viewer refreshes to show only those messages that you have allowed through the filters.

Note that in addition to filtering according to the categories, you can also display only those messages that contain a specific text string. To do so, click **Message** under **Set up Filters**, type the specific text for which you want to show messages, and select the check box for **Message** under **Activate Filters**.

Steps for tracing a business object flow

Flow tracing a business object, or access flow, allows you to track its progress throughout each of the processing points in its life cycle. Using Log Viewer, you can follow the progress by checking the trace messages that display. Each business object has an flow initiator ID associated with it for this purpose. If you sort the Log Viewer display by flow initiator ID and date/time, the trace messages for the business object are grouped together so you can easily follow its status. Sort by ascending or descending order to see a historical perspective or the latest event displayed first.

Note: Flow tracing is performed only when the business object is within the domain of InterChange Server Express, that is, from the connector controller of the source application to the connector controller of the destination application. Business object flow tracing is not performed when the business object is processed by connector agents or applications.

Perform the following steps to trace a business object flow:

1. Set the trace log file destination, if necessary (see “Configuring logging and tracing” on page 110).
2. Select the originating triggering business object (not the generic business object) to trace and set its trace level (see “Steps for configuring tracing levels for InterChange Server Express, business objects, and collaborations” on page 113).
3. Send an event from the source connector to the destination connector.
4. Open Log Viewer (see “Steps for starting Log Viewer” on page 116).
5. Set the display preferences to view the flow tracing (see “Overview of Log Viewer display options” on page 118).
6. Click any error message button in the **MsgID** column to view the text of the message.
7. Click any of the business object name buttons in the **BusObj** column to view the data contained within the business object.

This action uses the Business Object Viewer, which allows you to save the data to a separate file. The file can then be read by either the Mapping tool or the Test Connector.

Steps for controlling server memory usage

OutOfMemoryExceptions from the JVM can halt InterChange Server Express. To reduce such occurrences, a “memory checker” feature in InterChange Server Express can be used to pause the connectors when a predetermined level of memory usage is reached, ensuring that no new events are delivered to the connector, but still allowing it to process pending service call requests from InterChange Server Express. This can reduce memory usage. The parameters for the memory checking feature can also control the speed of event delivery to InterChange Server Express from all connectors, providing finer granularity of flow control.

Note: The memory checker feature can also *negatively* affect the performance of InterChange Server Express server, and should be used with caution. First try other methods of flow control, and use the memory checker only as the last solution for memory problems.

Perform the following steps to implement memory checking:

1. Add the following parameter to the InterChange Server Express startup script file, as a Java property parameter (-D option):

```
CW_MEMORY_MAX
```

This value determines the maximum heap memory InterChange Server Express is allowed to use. It should be the same value, in megabytes, as the JVM maximum heap size specified in the startup scripts of InterChange Server Express (-mx parameter). For example:

```
-DCW_MEMORY_MAX=512m
```

2. Use System Manager to customize server-memory usage by doing the following:

- In the InterChange Server Component Management view, right click the server instance and click **Edit Configuration**.
- Click the **Misc** tab and edit the values under **Server Memory**, as show in Figure 59:

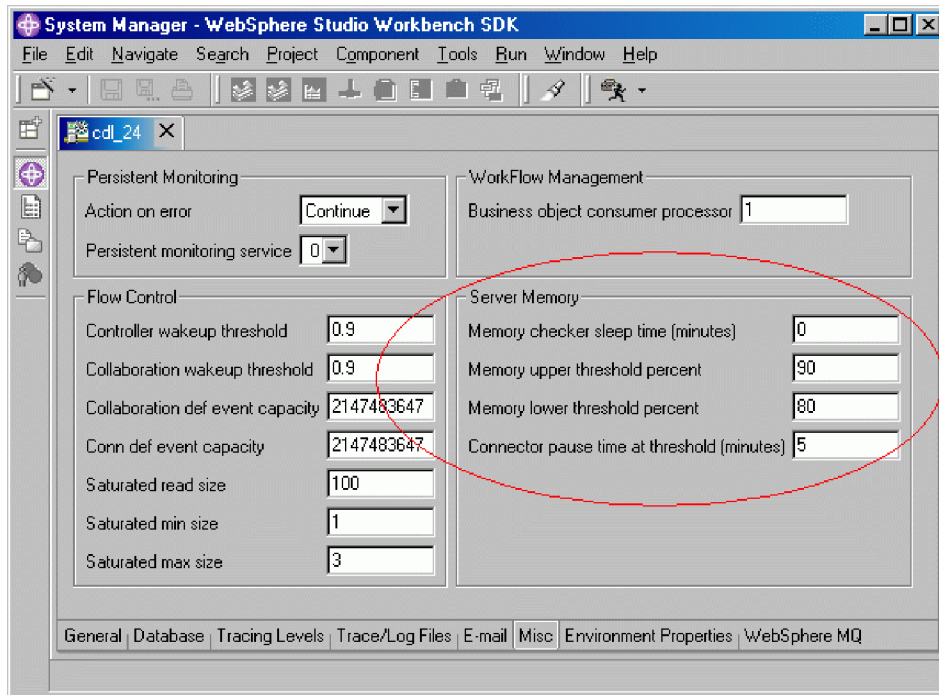


Figure 59. Edit Configuration tool, Misc Tab

Maintaining event archives

Some connectors use event archives at the application. If a connector archives events, you can use the archive to troubleshoot problems with the connector. The application administrator probably needs to flush the archive periodically so that it does not grow too large.

An event archive can provide a useful backup role. If a collaboration stops running, its subscriptions are canceled. When the collaboration is inactive, the connector continues to retrieve the associated events, but finds no active subscriptions, so it archives the events. When the collaboration restarts, you can examine the event archive and move archived events back to the event notification mechanism for reprocessing.

If a connector archives events, information on administration issues is available in that connector's documentation.

Attention: If you use WebSphere MQ for messaging, do not back up the MQ queues when performing regular backup procedures. The data in these queues is dynamic and, therefore, represents only in-progress transactions.

Managing Work-in Progress connections

Setting a sufficient minimum number of Work-in-Progress (WIP) connections helps eliminate the possibility of a server crash. A server may crash if it is attempting a reboot at the same time all connectors are trying to do WIP work and all collaboration threads are being committed to the WIP queue.

You should configure enough database WIP connections to accommodate potential run-time-driven work-load. Add the following property to the xml configuration file, `InterchangeSystem.cfg`:

```
MIN_CONNECTIONS = n
```

where

n = *sum_of_all_max_threads_for_collabs_used_simultaneously_plus_num_connectors*

sum_of_all_max_threads_for_collabs_used_simultaneously is the sum of the Maximum number of concurrent events settings of all the collaboration objects.

Managing failures

Managing failures in the InterChange Server Express system consists of using troubleshooting resources to resolve problems. Critical errors that may cause events to fail can occur with a system component such as a connector or collaboration, or a third-party component such as an integrated application.

When an error causes an event to fail, the InterChange Server Express system has built-in capabilities that let you resolve the problems. The system can be set up to pause collaborations within InterChange Server Express if a failure occurs. This section covers the following topics:

“Failure recovery for service calls” on page 123

“Strategies for InterChange Server Express recovery” on page 124

“Critical errors” on page 126

“Lost connection to application” on page 127

“Unknown connector agent status” on page 127

“Database connection failures” on page 128

“Flow failures” on page 128

Failure recovery for service calls

To avoid sending duplicate events to the destination application, you may want to prevent a recovery from automatically resubmitting all service calls that were in transit when a failure occurred. Before the server failure, you can configure a nontransactional collaboration to persist any service call event in the In-Transit state when a failure and recovery occurs. When InterChange Server Express recovers, the service call events remain in the In-Transit state, and you can use Flow Manager or Failed Event Manager to examine individual failed events and control when (or if) they are resubmitted.

To configure a collaboration to persist a failed service call in-Transit state, go to System Manager and select the **Persist Service Call In Transit State** check box in the Collaboration General Properties window.

Note: If the collaboration is part of a collaboration group, setting this property implicitly applies the same property setting to all other collaborations in the group. Do not set one collaboration for **Persist Service Call In Transit State** unless you want all other collaborations in that collaboration group to have that setting as well. For transactional collaborations, it is recommended that you not set this property; clear the **Persist Call In Transit State** check box.

Strategies for InterChange Server Express recovery

If InterChange Server Express fails when processing events, all the events currently in the Work-in-Progress (WIP) queue must be either recovered or otherwise processed when the server reboots. Potentially, because of memory requirements, the recovery of the WIP events can slow or even halt the server reboot. The InterChange Server Express product provides two features—deferred recovery and asynchronous recovery—for improving the time it takes the server to reboot and for making the server available for other work before all events have been recovered.

Flow control and storing business object keys as part of the WIP data assist in the efficiency of deferred recovery and asynchronous recovery. Both features reduce the amount of memory needed during an InterChange Server Express recovery, and therefore, can decrease significantly the amount of time necessary for InterChange Server Express to reboot during a recovery.

Storing business object keys as part of the WIP data means that during recovery, the business object key is retrieved without deserializing the business object, avoiding an MQ or a database round-trip. Flow control is a service that allows you to configure either system-wide or component-level queue depth parameters in to control the memory demands on InterChange Server Express. For more information about configuring flow control, see “Steps for configuring system-wide flow control” on page 40.

Steps for deferred recovery of collaboration events

In deferred recovery, recovery of a collaboration’s WIP events is deferred until after the server has rebooted, thereby saving the memory usage associated with those events.

After the server has rebooted, you can resubmit the events manually. Note the following recommendations:

- To avoid the possibility of data corruption due to event sequencing, resubmit the deferred events before you process any new events for that collaboration.
- Resubmit the deferred events in their original sequence.

You establish deferred recovery by setting the RECOVERY_MODE property of a collaboration object.

Note: Use of deferred recovery compromises the sequencing of events, which may cause data integrity problems. Use this feature only if sequencing is not important to you.

The RECOVERY_MODE property has two settings, which do the following when a server failure and reboot occurs:

- **Deferred**
Events that were in the Working state before the server failure are changed to the Deferred state. No events for this collaboration are recovered until you resubmit them manually.
- **Always**
Events that were in the Working state before the server failure are recovered. Events in the Deferred state remain deferred until you resubmit them.

The default setting is Always.

Note: Changing the collaboration recovery mode value from **Deferred** to **Always** does not recover deferred events, nor does it change existing deferred events to the Working state. Events in the deferred recovery state remain deferred until you manually resubmit them.

Perform the following steps to set the collaboration Recovery mode value:

1. From System Manager, right-click a collaboration object and click **Properties**. The Properties dialog box appears (see Figure 60).
2. In the Collaboration Properties dialog box, click the **Collaboration General Properties** tab. The following dialog box opens:

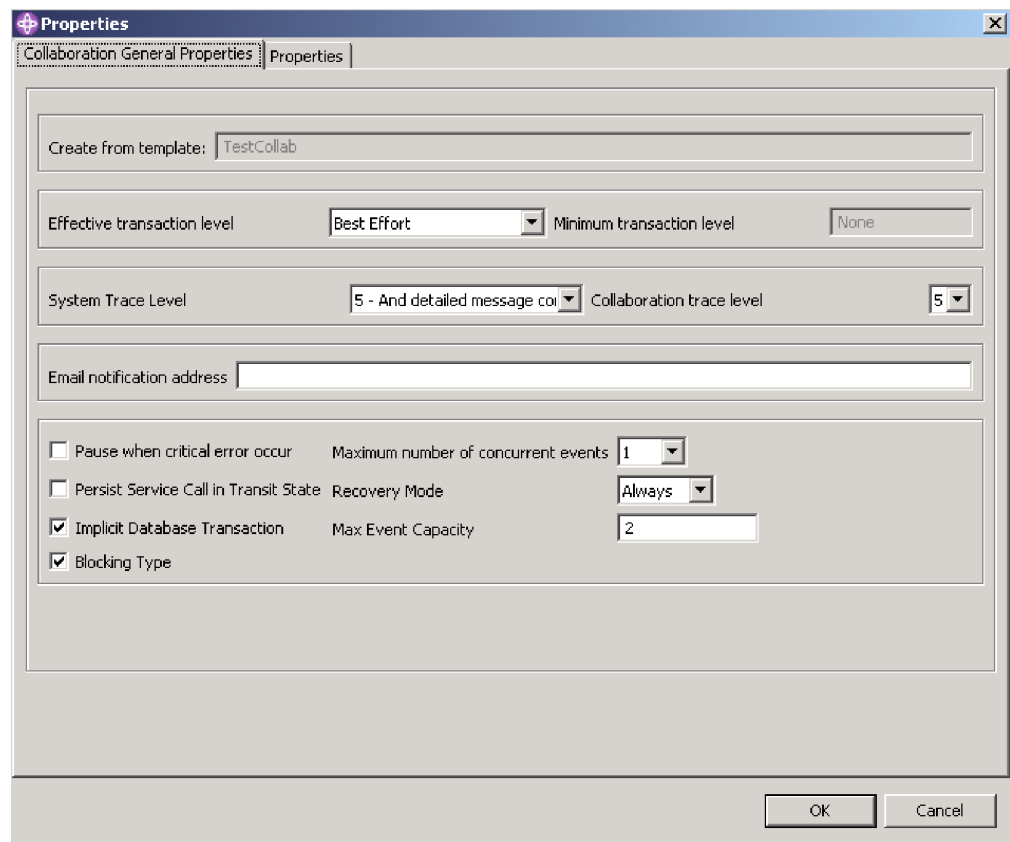


Figure 60. Properties dialog box, Collaboration General Properties tab

3. In the **Recovery Mode** list, select one of the following:
 - **Always**
The collaboration recovers all WIP events whose state is Working and that it owns at the time of server boot.
 - **Deferred**

The collaboration changes the WIP event to the Deferred recovery state. You must process those events at a later time using the Flow Manager or Failed Event Manager. For more information, see “Working with failed events” on page 129.

Asynchronous recovery

InterChange Server Express does not wait for collaborations and connectors to recover before it completes the boot process; collaborations and connectors are allowed to recover asynchronously after InterChange Server Express has booted. This makes it possible to use troubleshooting tools -- such as System Monitor, Failed Event Manager, and Flow Manager -- when the connectors and collaborations are recovering.

Critical errors

Critical errors in the InterChange Server Express system can cause problems in your run-time environment. A critical error as defined in the InterChange Server Express system can be generated by one of the following situations:

- Application time-out
- Inability to log on to the application
- Unknown connector agent status

By default, a collaboration continues processing subsequent initiators after a flow has failed. However, a collaboration’s behavior can be configured to pause automatically when a critical error occurs that might cause flows to fail. Configuring a collaboration in this way eliminates the possibility of the next flow failing for the same reason by not processing any more initiators after a flow fails. This is critical if the sequence in which initiators are processed needs to be maintained. If the collaboration pauses, the order in which initiators arrived to the server is maintained. At this point, you can fix the critical error, resolve the failed flow, then restart the collaboration. If collaborations do not depend on an initiator that is associated with a failed flow, you can resume the collaboration and resolve the failed flow at a later time. See “Flow failures” on page 128 for more information on submitting failed events.

To configure a collaboration object to pause after a critical error occurs, select the **Pause when critical error occurs** check box in the Collaboration General Properties tab of the Properties dialog box.

If this value is set, the collaboration pauses when a critical error occurs and remains paused until either of the following occurs:

- The connector agent reboots, automatically notifying the collaboration to resume processing, or
- You manually restart the collaboration, using Flow Manager or Failed Event Manager.

If you do not configure the collaboration to pause when a critical error occurs, the following situation might happen:

Two initiators, E1 and E2, are waiting to be processed by a collaboration. E1 creates a new customer and E2 updates E1. Because E2 updates E1, E1 must process before E2. If a critical error occurs when a collaboration is processing E1 and E1 fails as a result, then E1 is moved to the resubmission queue. If you do not select the **Pause when critical error occurs** check box, the collaboration attempts to process E2. E2 fails because it relies on the successful processing of E1.

If the collaboration property `CONVERT_UPDATE` is set to true, then E2, which updates E1, becomes a create and creates the new customer with the updated data. Data in E1 is now old and should not be manually submitted because it overwrites data delivered by E2.

Lost connection to application

Collaborations that are running assume that connectors have live connections to their applications. If a connector's application becomes unavailable, the connector is unable to poll the application for events and to satisfy collaboration requests.

When an application is unavailable, a connector that polls that application for events generates an error at each polling attempt. If the connector determines that the connection with the application has been lost, the connector agent terminates and returns a status to the connector controller requesting that the connector controller also terminate.

If a collaboration sends a request to a connector when the connector is up but its application has failed, the request returns with a failure status to the collaboration. This happens only if the connector property `ControlStoreAndForwardMode` is set to false. The collaboration fails, logging one of the following messages: 17050, 17058, 17059, or 17060. If you receive such messages, check the status of the application.

Unknown connector agent status

The status of the connector agent is crucial to the InterChange Server Express system because it is a starting point for application events. A connector controller maintains the status of its connector agent and relays this information to System Manager.

The connector controller maintains the status of its connector agent by sending response requests to the connector agent at 15-second intervals.

If the connector agent does not respond after three consecutive checks, its status is assumed to be unknown. An unknown connector agent status might mean that the connector agent has failed or if the connector agent is installed across the network, the network connection might have failed.

Setting the `ControllerStoreAndForwardMode` property for the connector to true makes the connector controller wait for the connector agent to start before delivering pending events. Setting this property to false makes the connector controller fail collaboration requests. The failed collaboration requests are moved to the resubmission queue and can be resubmitted using Flow Manager. See "Flow failures" on page 128 for more information.

Note: Collaborations relying on an agent whose status is unknown can stop responding if the `ControllerStoreAndForwardMode` property is set to true. Collaborations stop responding until the connector agent is restarted and replies to the collaboration requests.

When the **Pause when critical error occurs** check box is selected for the associated collaborations, the collaborations bound to this connector pause upon receiving the unknown status of the connector agent. An error message is logged and e-mail is sent if the e-mail connector is configured.

Database connection failures

When InterChange Server Express needs a database connection for one of its services but finds that the maximum number of connections are already in use, the server tries to free a connection that is idle. If the server is unsuccessful, the connection attempt fails and InterChange Server Express logs error 5010: Unable to find an available connection in the cache. The maximum number of connections *max-connections-value* has been reached.

If you set a constraint on the number of InterChange Server Express connections by setting the `MAX_CONNECTIONS` parameter, you should monitor error 5010 messages because a connection failure can have undesirable consequences. For example, when InterChange Server Express cannot obtain a connection for its event management service, it stops running. By default, this constraint is set to an unlimited number of connections.

Connection failures indicate that the maximum number of allocated connections is insufficient to meet the run-time work load. If you cannot allocate more connections to InterChange Server Express in the current database, consider partitioning its workload across multiple databases.

Flow failures

At times, the InterChange Server Express system or its associated applications may fail. Successfully processing flows that carry data through the InterChange Server Express system is critical, so in a run-time environment it is critical to maintain data consistency. System failures such as system errors, data errors, and critical errors can cause flows to fail to process. The InterChange Server Express system has built-in capabilities that allow you to process system failures.

A system configuration, object definition, application-specific, or data consistency error can cause a flow to fail when the InterChange Server Express system is processing that flow. Improperly functioning InterChange Server Express components, such as business object mapping failures or the unavailability of a connector, can generate system errors, which cause flows to fail. Data inconsistencies, such as an isolation violation of application data during execution of a collaboration, generate data errors, which also cause flows to fail.

If an error occurs when a connector controller or a collaboration is processing a flow, the flow fails and is moved to the event resubmission queue. From here, you have the following choices:

- Submit the event with the original business object
- Submit the event using the latest business object
- Discard the event from the system

For instructions on resolving failed flows, see “Working with failed events” on page 129..

Overview of failed transactional collaborations

System and data errors can cause a transactional collaboration to fail. When one of these errors occurs, the collaboration attempts a rollback. If the rollback of a collaboration’s compensation steps fails, the collaboration is in an “in-doubt” state. If an error occurs during run-time recovery, the collaboration is put into a list of failed transactional collaborations owned by the corresponding collaboration. A *failed transactional collaboration* is a collaboration whose compensation steps failed to roll back.

After a transactional collaboration fails, you must resolve it. You can process a failed transactional collaboration by using Flow Manager or Failed Event Manager. For instructions on resolving failed transactional collaborations, see “Working with failed events” on page 129..

Steps for preventing failed transactional collaborations from pausing

The default behavior for a failed transactional collaboration is to pause. You can prevent failed transactional collaborations from pausing by adding a property called PAUSE_ON_COMPENSATION_FAILURE to the collaboration template and changing the setting from TRUE (default) to FALSE.

Perform the following steps to add the new property and change the setting to FALSE:

1. In System Manager, double-click the collaboration template that failed. Process Designer Express opens.
2. Double-click the Definitions icon under Site Wrapper. The Template Definitions window appears in the right-hand frame.
3. Click the **Properties** tab.
4. Click **Add**. A Name dialog box appears.
5. Type PAUSE_ON_COMPENSATION_FAILURE in the Name field, then click **OK**. The PAUSE_ON_COMPENSATION_FAILURE property appears under the General Properties node in the left-hand pane.
6. With PAUSE_ON_COMPENSATION_FAILURE selected in the left-hand pane, select **Boolean** from the **Property Type** list.
7. Under **Value**, clear the **IsDefaultVal** check box in the **True** row, and select the **IsDefaultVal** check box in the **False** row.
8. Click **Apply**.
9. Close Process Designer Express.

Working with failed events

You can use two tools in IBM WebSphere InterChange Server Express 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 InterChange Server Express product. This section contains the following:

“Using Failed Event Manager”

“Using Flow Manager” on page 136

Using Failed Event Manager

Failed Event Manager allows you to view and manage failed events from the Web.

Failed Event Manager 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 Express Installation Guide*. Failed Event Manager works with the following versions of WebSphere Application Server, WebSphere Application Server Express, and Tomcat:

- WebSphere Application Server versions 5.0.2 or 5.1
- WebSphere Application Server Express versions 5.0.2 or 5.1

If you selected an Administrative Toolset installation of WebSphere Business Integration Server Express, Failed Event Manager is automatically installed and configured if WebSphere Application Server or WebSphere Application Server Express 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 Express Installation Guide*). Default roles mentioned above 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”

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 creating custom users and roles for Failed Event Manager with Tomcat”
- “Steps for logging on to Failed Event Manager” on page 132
- “Steps for viewing failed events” on page 133
- “Steps for processing failed events in Failed Event Manager” on page 135
- “Steps for checking your access rights in Failed Event Manager” on page 136
-

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 Express, 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**.

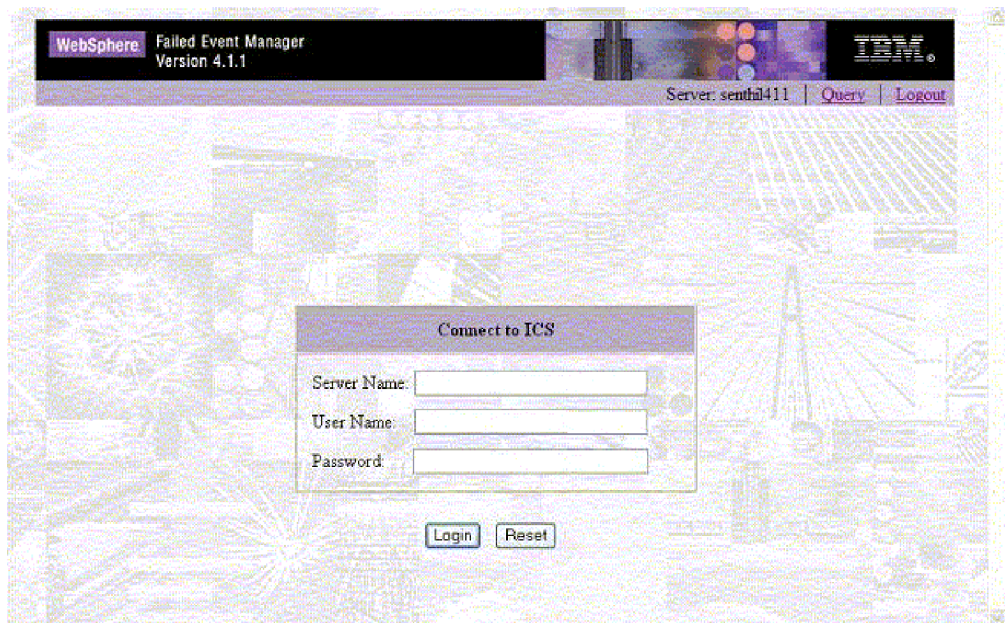


Figure 61. Failed Event Manager Connect screen

After you log on, the Query dialog box appears, as shown in Figure 62.

Query

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

Submit Reset

[Check your access rights](#)

Figure 62. 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 many events you want to be displayed on each page of Failed Event Manager.
2. Click **Submit**.

A Table of Failed Events appears, as shown in Figure 63 on page 134.

WebSphere Failed Event Manager Version 4.1.1 Server: sentsh411 Query Logout

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

Submit Refresh & Submit Cancel Waiting Delete

Figure 63. 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 64.

WebSphere Failed Event Manager Version 4.1.1 Server: sentsh411 Query Logout

Table of Failed Events

Now showing failed events 1 to 5 of 20.

Status	PointOfFailure	Connector	BusinessObject	Time	Message	Key[s]	Details
<input type="checkbox"/> Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:36:17.14	Error11085 Collaboration exception	Id=564	View
<input type="checkbox"/> Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:36:46.933	Error11085 Collaboration exception	Id=23	View
<input type="checkbox"/> Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:37:16.887	Error11085 Collaboration exception	Id=24	View
<input type="checkbox"/> Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:36:51.37	Error11085 Collaboration exception	Id=25	View
<input type="checkbox"/> Failed	AutoTestTransCollabObj	AutoTestConnector	SimpleBO.Create	2004-01-16 14:37:22.103	Error11085 Collaboration exception	Id=26	View

Submit Refresh & Submit Delete

Figure 64. Failed Event Manager Table of Failed Events without rights to view business objects

- 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 65 on page 135.



Figure 65. 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 66.

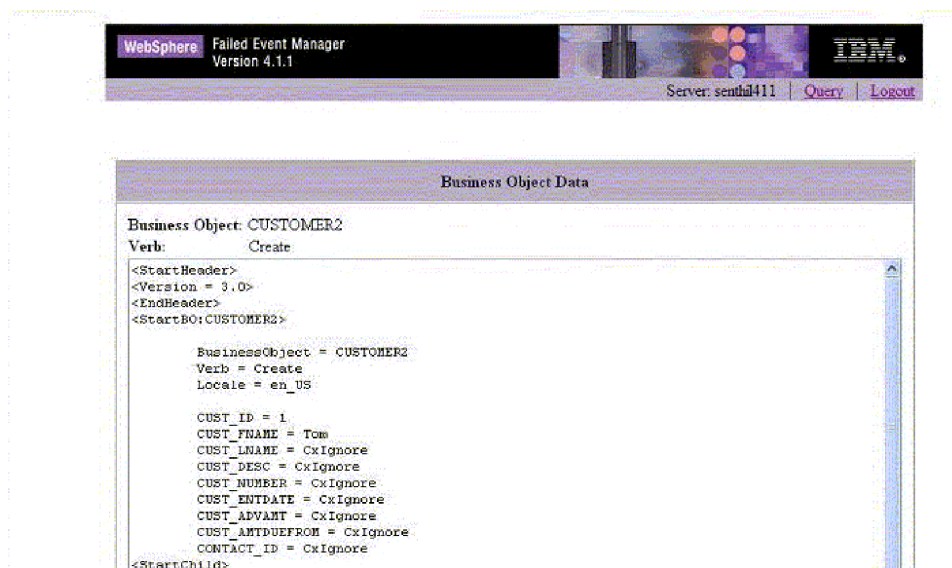


Figure 66. Failed Event Manager Business Object Data dialog box

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 132).
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 InterChange Server Express 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. For background information on failed flows, see “Flow failures” on page 128.

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 138

“Steps for viewing details for unresolved flows” on page 141

“Steps for processing unresolved flows” on page 143

Steps for starting Flow Manager

Perform the following steps to start Flow Manager:

1. Click **Start > Programs > IBM WebSphere Business Integration Express > Toolset Express > Administrative > Flow Manager**. The Connect to InterChange Server dialog box appears (see Figure 67 on page 137).

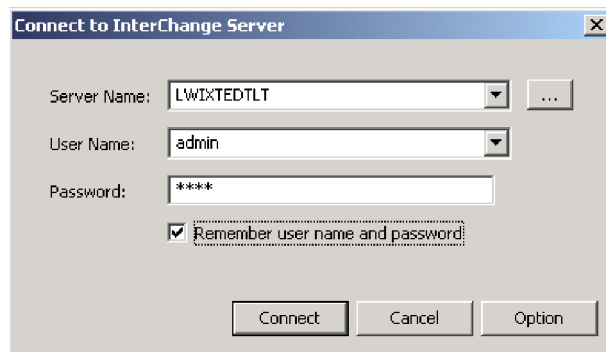


Figure 67. Connect to InterChange Server dialog box

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.

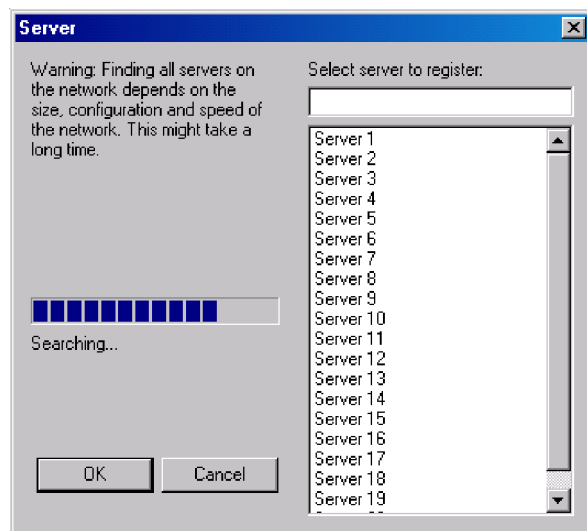


Figure 68. Server dialog box

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.




Figure 69. Options dialog box

After you click **Connect** in the Connect to InterChange Server dialog box, Flow Manager opens (see Figure 70).




Figure 70. 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 136.
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 139
 - “Date & Time tab” on page 140
 - “Error Text & Business Object Attributes tab” on page 140
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 71 on page 139).




Figure 71. 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 72).




Figure 72. 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 73).

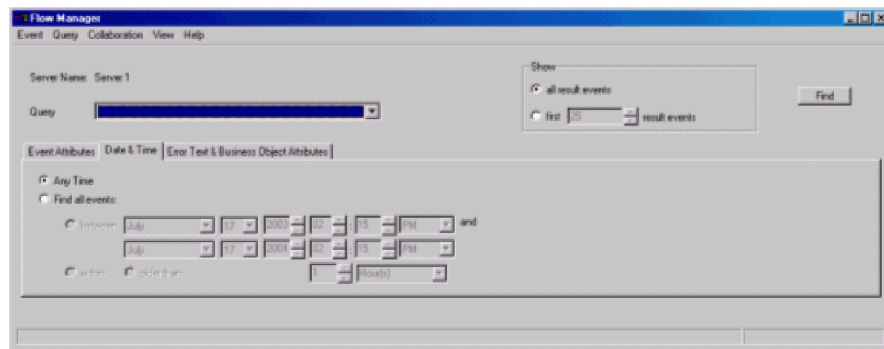


Figure 73. Date & Time tab

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 74 on page 141).




Figure 74. 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 71 on page 139).

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...]
```

- 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 75).

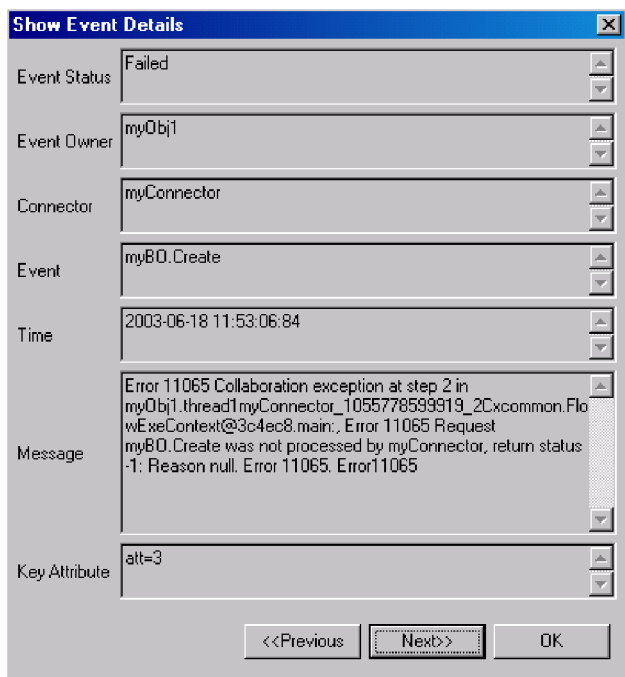


Figure 75. Flow Manager Show Event Details dialog box

- 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 76).

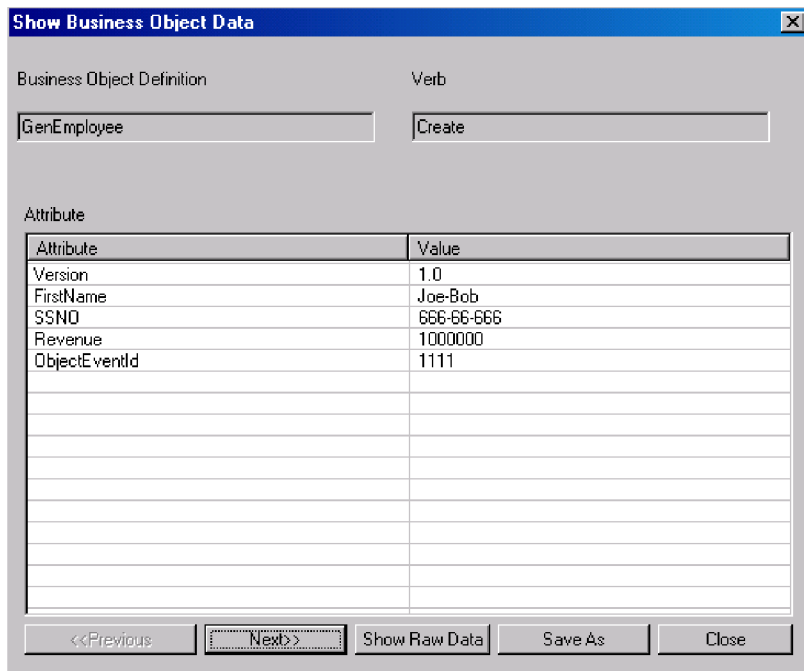


Figure 76. Show Business Object Data dialog box

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**.

Resolving collaboration deadlocks

A deadlock is a situation where two or more processes cannot proceed because each is waiting for the other processes to proceed. Deadlocks are an undesirable side effect of the concurrency control provided by event isolation within a collaboration. See the *Collaboration Development Guide* for more information on event isolation.

Figure 77 on page 144 illustrates a deadlock between two active collaboration groups resulting from the following sequence of events:

1. At time T1, Collaboration A1 receives an event, E1, then makes a service call to Collaboration B2 and sends a child business object for E1. Collaboration A1 waits for the service call to complete.

2. At time T2, Collaboration B1 receives an event, E2, then makes a service call to Collaboration A2 and sends a child business object for E2. Collaboration A2 waits for the service call to complete.
3. At time T3, Collaboration B2 is waiting for Collaboration B1, because B2 and B1 have the same port binding and the event from B1 was delivered before the event for B2 arrived.
4. At time T4, Collaboration A2 is waiting for Collaboration A1, because A2 and A1 have the same port bindings and the event from A1 was delivered before the event for B1 arrived.

At this point, all collaborations cannot move forward.

Note: A port binding consists of the business object type and connector name. See the *Collaboration Development Guide* for more information on port bindings.




Figure 77. Deadlock between collaboration groups

This section covers the following topics:

“Steps for detecting a collaboration deadlock” on page 144

“Steps for detecting group collaboration deadlocks” on page 145

“Steps for fixing a collaboration deadlock” on page 145

Steps for detecting a collaboration deadlock

By default, the InterChange Server Express system performs deadlock detection automatically when you start InterChange Server Express. However, when deadlock detection is performed, startup of InterChange Server Express can be delayed if a collaboration group contains many collaboration objects, because InterChange Server Express must traverse all the collaboration objects in the group to determine whether there is a deadlock in the group. This can cause slow startup even when no deadlock exists.

You can configure the InterChange Server Express system to not perform deadlock detection. If you do so, the system starts collaboration groups without first checking for deadlocks. This can make it possible for InterChange Server Express

to boot more quickly. However, if deadlock detection is not performed and a deadlock exists, an event that is later sent to a collaboration may fail.

System Manager does not provide the ability to set the `DEADLOCK_DETECTOR_CHECK` configuration parameter. Instead, to set this configuration parameter, you must edit the `InterchangeSystem.cfg` file and change the parameter's value in this file.

Perform the following steps to configure the InterChange Server Express system for deadlock detection:

1. Open the `InterchangeSystem.cfg` file. The following lines define the `DEADLOCK_DETECTOR_CHECK` parameter in the file:

```
<tns:name>DEADLOCK_DETECTOR_CHECK</tns:name>  
<tns:value xml:space="preserve">false</tns:name>
```
2. If you want deadlock detection performed, change the false value to true.
3. If you do not want deadlock detection performed, change the value to false.

Steps for detecting group collaboration deadlocks

Perform one of the following steps to check for a group collaboration deadlock:

- In System Manager, right-click the running group collaboration and click **Diagnostics**.

A window appears with the following message:

The following diagnostic tests were run on this collaboration:

This message is followed by one of the following results:

- Deadlock detection - Fail
- Deadlock detection - Pass
- No diagnostic tests supported at this time

Note: This result appears only if you run Diagnostics on a non-group collaboration.

- Check the `InterchangeSystem.log` file for the following error at the time the hanging collaborations were started:
Error 11135: Activation of collaboration *collaboration_name* group could cause a potential deadlock with one or more existing collaboration groups, and is therefore disallowed.

This error warns only of a potential deadlock situation. The informational messages preceding error 11135 identify the active collaboration groups that potentially enter into a deadlock.

Steps for fixing a collaboration deadlock

Perform the following steps to fix a collaboration deadlock:

1. Gracefully shut down all other collaborations.
2. Shut down the server immediately. See “Shutting down InterChange Server Express” on page 35 for more information on shutting down the InterChange Server Express system.
3. Restart the system. Upon system restart, a stopped collaboration that caused the deadlock automatically starts and resubscribes to all of the business objects it supports. The business objects that caused the collaborations to enter into a deadlock are redelivered. Because deadlocks are timing-dependent, the collaborations are prevented from entering into another deadlock. It is unlikely

that you have the same server load and isolation sequencing that you had when your system encountered the deadlock.

4. After restarting the system, shut down the collaborations involved and rebind the ports so that this does not occur again.

Steps for preventing collaboration deadlocks

You can prevent collaboration deadlocks by configuring the deadlock retry settings in the Database tab of server configuration screen in System Manager.

Perform the following steps to configure the deadlock retry mechanism:

1. In System Manager, right-click the server in the InterChange Server Component Management view and click **Edit Configuration**. The upper-right section of the System Manager window becomes a tool in which you can edit the InterchangeSystem.cfg file.
2. Click the **Database** tab. A dialog box appears in the upper-right section of the System Manager window in which you can enter the parameters necessary for configuring database configuration at the system level (see Figure 78).

The screenshot shows the 'Database Configuration' dialog box for the server 'LWIXTEDTLT'. The 'Database Connectivity' section is expanded, showing the following settings: Database Driver: MS SQL Server; Max Connections: 50; Max pools: 10; Idle Timeout: 2 Min; Max deadlock retry: 5; Deadlock retry interval: 20 Second(s). Below this, the 'Event Management', 'Transactions', and 'Repository' sections are visible, each with fields for Host name, Database, Login, Password, Max Connections, and Port number. The 'Flow Monitoring' section is also visible but has empty fields. The 'Database' tab is selected at the bottom of the dialog.

Figure 78. Edit Configuration screen, Database tab

3. In the **Max database retry** field, type a number that represents the maximum number of retries you want the server to perform if a deadlock occurs.
4. In the **Deadlock retry interval** field, type a number that represents the number of seconds you want the system to wait before retrying.

Managing IBM WebSphere MQ

This section contains information about resolving problems related to IBM WebSphere MQ. The following topics are covered:

“Queue configuration errors” on page 147

“Failure to create new instances of MQQueueManager” on page 147

“Failure recovery and the queue manager” on page 147

“Memory checker segmentation fault” on page 147

Queue configuration errors

IBM WebSphere MQ may report errors if the configuration settings for the queues do not allow for a sufficient number of messages, a sufficient maximum length for messages, or a sufficient size for log files. For information about adjusting these configuration settings, see *WebSphere Business Integration Server Express Installation Guide*.

Failure to create new instances of MQQueueManager

On the AIX platform, the use of the bindings transport (C via JNI) in JDK 1.3.1_06 requires additional configuration to avoid a conflict with shared memory segments. Without this configuration, the system might not create a new instance of MQQueueManager (reason code 2059). The failure can be accompanied by the generation of FDCs in the `/var/mqm/errors` directory.

To perform the configuration, use an additional parameter for the `mqs.ini` file. The parameter is `IPCCBaseAddress` and is set on a per queue manager basis. By default this parameter is set to the value 8, but it is recommended that you set it to 12. The following is an example of an altered QueueManager stanza in an `mqs.ini` file:

```
QueueManager:  
  Name=MQJavaTest  
  Prefix=/var/mqm  
  IPCCBaseAddress=12
```

Failure recovery and the queue manager

After a full system failure and recovery, the queue manager might fail to start, displaying this message:

```
AMQ7017 Log not available.
```

This typically indicates that the log file is missing or damaged, or that the log path to the queue manager is inaccessible.

To remedy the problem, create a dummy queue manager and copy its header log file and transaction log files over to the actual queue manager. This allows you to restart the queue manager.

Memory checker segmentation fault

The memory checker thread can get a segmentation fault when JMS or other modules are used in WebSphere MQ binding mode.

By default, Java 1.3.1 runs using eight segments. (Eight is the maximum value allowed; each segment is 256 MB). You can reduce the number of segments Java

uses with the LDR_CNTRL environment variable. For example, if you want Java to use only three segments, set LDR_CNTRL to set the following value prior to starting Java:

```
export LDR_CNTRL=MAXDATA=0x30000000
```

If your native methods use a large amount of shared memory (for example, with functions like `shmat()` and `mmap()`), you might need to reduce the number of segments used by Java. You might also find that you need to reduce the number of segments when concurrently running Java with other applications that use shared memory, particularly those that explicitly specify the address where a shared memory segment is attached (rather than letting AIX choose the address).

Steps for managing database connection pools

At run time, it takes time for processes to establish new connections to the database. You can minimize that time by establishing database connection pools for use by your collaboration and map processes.

Use of database connection pools requires establishing pool names in collaboration template and map code, in addition to creating the database connection pools in System Manager. For instructions on performing those tasks, and additional information about setting up database connection pools, see the *System Implementation Guide*.

Perform the following steps to manage how the database connection pool resources are allocated to your collaboration and map processes (if database connection pools have been set up on your system):

1. Revise the number of connection processes that can be used by a specific database connection manager and its pools.
2. Revise the database connection pool assignments.

For details about these tasks, see the *System Implementation Guide*.

Steps for managing run-time properties

The components of the WebSphere InterChange Server Express 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 InterChange Server Express system read this file as part of their initialization process:

- An adapter
- InterChange Server Express
- `repos_copy` utility
- Many InterChange Server Express tools

This `CWSharedEnv` file contains initialization information for the communication and run-time software that the InterChange Server Express 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.

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.”

Managing the Object Request Broker

This section contains information about resolving problems related to the IBM Java Object Request Broker (ORB), which handles communication between InterChange Server Express and several of its components. This section provides information on the following topics:

- “Object Request Broker installation”
- “Steps for customizing the Object Request Broker”
- “Steps for changing the location of the Object Request Broker” on page 151
- “Using the IBM Transient Name Server” on page 152

Object Request Broker installation

InterChange Server Express 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 WebSphere InterChange Server Express Installer installs automatically.

Steps for customizing the Object Request Broker

Perform the following steps to customize the Object Request Broker:

1. Refer to Table 23 on page 150 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, InterChange Server Express instances, the `repos_copy` utility, and the InterChange Server Express tools.

Note: For more information on the CWSHaredEnv file, see “Steps for managing run-time properties” on page 148.

As Table 23 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 InterChange Server Express configuration file (InterchangeSystem.cfg) and the adapter local configuration file can contain a CORBA section. When you specify the configuration parameter in the CORBA configuration file, you change it for all ORB-related tasks that the ORB server performs.

Important: The configuration files are in XML format. Do *not* modify these files unless you use an XML editor or are very familiar with XML format!

For example, to specify the maximum number of threads, you can take any of the following actions:

- Add the IBM ORB property to the ORB_PROPERTY variable in the CWSharedEnv file:

```
ORB_PROPERTY=-DORBNamingProvider=CosNaming
-Dorg.omg.CORBA.ORBClass=com.ibm.CORBA.iiop.ORB
-Dorg.omg.CORBA.ORBInitialPort=%ORB_PORT%
-Dorg.omg.CORBA.ORBInitialHost=%ORB_HOST%
-Dcom.ibm.CORBA.Debug.Output=null
-Dcom.ibm.CORBA.ThreadPool.MaximumSize=100
```

- Specify the IBM ORB property on the command line when you start up the component:

```
start_server.....-Dcom.ibm.CORBA.ThreadPool.MaximumSize=100
```

- Add the OAtHreadMax configuration parameter to the CORBA section of the configuration file:

```
[CORBA]
OAtHreadMax=100
```

Table 23. IBM Java ORB properties that can be customized

IBM ORB property	Configuration parameter	Description
com.ibm.CORBA.ListenerPort	OApport	Port number on which the ORB server (within InterChange Server Express) listens for incoming requests.
com.ibm.CORBA.LocalHost	OAIpAddr	IP address or host name of the machine on which the ORB server (within InterChange Server Express) is running.

Table 23. IBM Java ORB properties that can be customized (continued)

IBM ORB property	Configuration parameter	Description
com.ibm.CORBA.ThreadPool.MaximumSize	OAThreadMax	Maximum number of threads that the connection manager can create. The default value (zero) indicates that no size restriction exists.
com.ibm.CORBA.ThreadPool.InactivityTimeout	OAThreadMaxIdle	The time (in seconds) before an idle thread 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 24 for the default ORB-location information.
2. Change this default information during installation. In the Naming Server screen, the WebSphere InterChange Server Express Installer prompts you for the IP address and port number for the InterChange Server Express instance. The installer saves this information in the appropriate variables of the product directory.

Table 24 also shows the variables within the CWSHaredEnv file that specify the ORB location.

Table 24. Location of the IBM Java ORB

ORB location	IBM ORB property	Default value	CWSHaredEnv variable
IP address	org.omg.CORBA.ORBInitialHost	Name of the local host	ORB_HOST
Port number	org.omg.CORBA.ORBInitialPort	14500	ORB_PORT

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 24 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 24 on page 151 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 149.

Using the IBM Transient Name Server

Using IBM Transient Naming Server (tnameserv) provides the naming service for the WebSphere Business Integration Express system. When a component of the WebSphere Business Integration Express system starts, it registers itself with the IBM Transient Naming Server. When the component needs access to another business-integration-system component, it uses the naming service to determine the information it needs to locate and start interacting with that component. For example, when an adapter needs to communicate with InterChange Server Express, it obtains the location of InterChange Server Express through the Transient Naming Server.

Note: The IBM Transient Naming Server is part of the IBM Java ORB. Therefore, it is installed automatically on the InterChange Server Express machine as part of the WebSphere InterChange Server Express installation process.

Identification of registered components

The WebSphere InterChange Server Express product provides the `CosNameServer_Dump` tool to list all valid InterChange Server Express 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 WebSphere Business Integration Express 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 InterChange Server Express 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 InterChange Server Express configuration file (InterchangeSystem.cfg).
2. Set the CosNamingPersistenceFile 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, InterChange Server Express updates the naming repository by copying the CORBA objects currently registered with the Transient Naming Server into the naming-repository file. When each adapter starts, it updates the naming repository with its information. If InterChange Server Express has not yet started when the adapter starts, the naming repository is updated whenever InterChange Server Express 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 InterChange Server Express or any started adapters.

Perform the following steps to turn off the Persistent Naming Server:

1. Edit the InterChange Server Express configuration file (InterchangeSystem.cfg).
2. Set the CosNamingPersistence configuration parameter (located in the CORBA section) to false.

Appendix A. IBM WebSphere MQ reference

This appendix describes some of the commands that are used at the DOS command prompt to administer IBM WebSphere MQ. A complete description of the WebSphere MQ commands can be found in the WebSphere MQ online documentation.

Command	What the command does
runmqsc or start_mq.bat	Starts the WebSphere MQ command interpreter. The remaining commands can be entered after this command is run.
endmqm -i queue.manager	Ends WebSphere MQ queues immediately. Must be run before deleting queues.
clearorclear_mq.bat	Clears WebSphere MQ queues.
dltmqm queue.manager	Deletes WebSphere MQ queues.
crtmqm queue.manager	Creates a WebSphere MQ queue manager.
configure_mq.bat <path to crossworlds_mq.tst>	Configures a WebSphere MQ queue manager.
define	Defines WebSphere MQ queues.

Appendix B. Requirements for restarting IBM WebSphere InterChange Server Express system components

This appendix describes the restart requirements for administrative tasks. The restart requirements for development or implementation tasks are described in the appropriate development guide or in the *System Implementation Guide*.

Use the following tables to decide if it is necessary to restart individual IBM WebSphere InterChange Server Express components, including InterChange Server Express. An entry labeled Dynamic in the Restart Requirements column means the component does not require a restart.

- “InterChange Server Express 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 Express restart requirements

Table 25 describes the restart requirements for InterChange Server Express.

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 25. InterChange Server Express 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 Express
Change InterChange Server Express name	Stop InterChange Server Express, re-create queues with new names, then start InterChange Server Express
Change InterChange Server Express password	Dynamic
Change repository database passwords	Dynamic
Add new class libraries (*.jar)	Add the name to the startup file, then restart InterChange Server Express
Change class libraries	Restart InterChange Server Express

Collaboration template restart requirements

Table 26 on page 158 describes the restart requirements for collaboration templates.

Note: These restart requirements are for administrative tasks only. For information on restart requirements for development or implementation tasks, see the appropriate development guide or the *System Implementation Guide*.

Table 26. 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 27 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 27. 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 28 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 28. Connector restart requirements

Action	Restart requirement
Add a new connector	Restart InterChange Server Express
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 Express 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 28. Connector restart requirements (continued)

Action	Restart requirement
Change transport protocol	Restart InterChange Server Express
Change concurrent event-triggered flows	Restart InterChange Server Express
Change poll start time	Restart connector
Change poll end time	Restart connector agent
Create new connector	Does not require a restart of InterChange Server Express
Recompile connector	Restart connector agent
Add 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 Express. Until InterChange Server Express is restarted, the Associated Maps screen is blank because the information in this screen reflects the InterChange Server Express run time only. Because this is a new connector with new supported business objects that are not yet reflected in InterChange Server Express run time, the Associated Maps screen has no running business objects for binding maps.
Make changes to the values of the following standard properties:	From System Manager, right-click the server under Server Instances, then select Refresh. The Refresh option stops the connector, temporarily caches its transient state, removes the instance from memory and replaces it with the new instance, then starts the new instance.
<ul style="list-style-type: none"> • DeliveryTransport • AllowAnonymousConnections • EbGateway • RemoteWebGatewayURL • ListenPort • CACertificateDirectory • ConcurrentEvents • JMSBrokerName • JMSFactoryClassName • JMSUserNama • JMSPassword • NumberOfConnections 	

Business object restart requirements

Table 29 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 29. 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 Express or Process Designer Express, is connected to InterChange Server Express 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 Express is strongly recommended.

Table 29. Business object restart requirements (continued)

Action	Restart requirement
Change key attribute.	Restart connector agent.
Change default value.	If the connector uses the "UseDefaults" connector configuration property, the connector agent must be restarted. Restarting InterChange Server Express 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/subobjects).	Stop InterChange Server Express.
	Clean queues.
	Restart InterChange Server Express.
	The maps using the object must be updated and recompiled.

Map restart requirements

Table 30 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 30. 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 Express restarted.
Change imported class libraries	Restart InterChange Server Express
	If there are signature changes, the maps must also change and must be recompiled.

Relationship restart requirements

Table 31 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 31. 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
CrossWorlds
DB2
DB2 Universal Database
Domino
Lotus
Lotus Notes
MQIntegrator
MQSeries
Tivoli
WebSphere

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

MMX, Pentium, and ProShare are trademarks or registered trademarks of Intel Corporation in the United States, other countries, or both.

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

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

System Manager includes software developed by the Eclipse Project (<http://www.eclipse.org/>).



IBM WebSphere Business Integration Server Express version 4.3 and IBM WebSphere Business Integration Express Plus version 4.3.

Index

A

- about logging 99
- administering
 - collaboration objects 52
 - connectors 42
 - InterChange Server Express 34
 - InterChange Server Express components 33
 - maps 58
 - relationships 60
- archive tables, backing up 86
- automatic and remote restart of connectors 48

B

- backing up
 - archive tables 86
 - collaboration class files 86
 - components 85
 - repository 85
- backup
 - planning 84
 - schedule 84
- bar display option, default monitors 12
- business object probe data file, deleting in System Monitor 22
- business objects
 - restart requirements 159
 - tracing 106

C

- collaboration class files, backing up 86
- collaboration deadlocks
 - deleting 144
 - detecting group collaboration deadlocks 145
 - fixing 145
 - preventing 146
 - resolving 143
- collaboration objects
 - administering 52
 - configuring concurrent event-triggered flow processing 57
 - configuring flow control 57
 - configuring run-time properties 55
 - configuring trace levels 114
 - log messages 100
 - operating 52
 - restart requirements 158
 - scheduling 80
 - setting general properties 55
 - starting, stopping, and pausing 54, 55
 - states 52
 - statistics 28
 - tracing 105
 - viewing states in System Manager 53
 - viewing states in System Monitor 52
- collaboration templates, restart requirements 157
- collaborations
 - deferred recovery of WIP events 124
 - preventing failed transactional from pausing 129
- concurrent event-triggered flows, configuring for collaboration objects 57

- config file 110
- configuring
 - collaboration object trace level 114
 - collaboration tracing 114
 - connector agent logging destination 112
 - e-mail notification 107
 - flow control 40, 51, 57
 - InterChange Server Express logging and tracing parameters 112
 - InterChange Server logging and trace destinations 110
 - logging 110
 - map tracing 115
 - persistent monitoring 23
 - stack trace 112
 - system-wide flow control 40
 - trace levels 113
 - tracing 110
- connecting to an InterChange Server Express instance 25
- connector agent
 - configuring logging destination 112
 - log messages 102
 - unknown status 127
- connectors
 - administering 42
 - automatic and remote restart 48
 - changing states in System Manager 45
 - configuring flow control 51
 - configuring tracing 114
 - enabling for MQ-triggered OAD 49
 - initializing 44
 - log messages 100
 - operating 42
 - restart requirements 158
 - restarting 48
 - restarting in System Manager 48
 - restarting in System Monitor 48
 - running, stopping, and pausing 44
 - scheduling 80
 - shutting down 47
 - shutting down in System Manager 48
 - shutting down in System Monitor 47
 - starting manually 45
 - states 42
 - statistics 30
 - tracing 105
 - viewing states in System Manager 43
 - viewing states in System Monitor 42
- creating schedules 81
- critical errors 126
- customizing the Object Request Broker (ORB) 149

D

- database connection failures 128
- database connection pools, managing 148
- database, changing password 36, 37
- deleting schedules 83
- disabling schedules 83
- disconnecting from InterChange Server Express 67
- displaying schedules 83

E

- e-mail notification
 - collaboration object level 109
 - configuring 107
 - connector level 109
 - system level 108
- Edit Configuration tool 110
- enabling schedules 83
- event archives, maintaining 122

F

- Failed Event Manager 129
 - checking access rights 136
 - creating custom users and roles with Tomcat 130
 - logging on 132
 - processing failed events 135
 - viewing failed events 133
- failed events, working with 129, 133, 135
- failed transactional collaborations 128
- failure recovery for service calls 123
- failures
 - database connection 128
 - flow 128
 - managing 123
- flow control
 - configuring for collaboration objects 57
 - configuring for connectors 51
 - configuring system-wide 40
- flow failures 128
- Flow Manager
 - finding unresolved flows 138
 - handling unresolved flows 143
 - managing queries 143
 - starting 136
 - viewing details of unresolved flows 141
- flow tracing 106

H

- historical data, setting the frequency for capturing in System Monitor 20
- historical statistics, deleting in System Monitor 22

I

- IBM Transient Name Server 152
- initializing a connector 44
- installing
 - Object Request Broker (ORB) 149
- InterChange Server Component Management view 25
- InterChange Server Express
 - administering 34
 - changing password 36
 - configuring logging and tracing parameters 112
 - connecting to an instance 25
 - customizing startup parameters 34
 - disconnecting from 67
 - graceful shutdown 35
 - immediate shutdown 36
 - logs 102
 - operating 34
 - password 36, 157
 - restart requirements 157
 - shutting down 35

- InterChange Server Express (*continued*)
 - starting 34
 - startup parameters 34
 - statistics 26
 - strategies for recovery 124
 - tracing services 105

J

- jobs, scheduling 78

L

- line display option, default monitors 11
- log file management 103
- log messages
 - viewing 115
- log viewer
 - filtering messages 120
 - setting preferences 116
 - starting 116
 - tracing business object flow 120
 - using for tracing 115
- logging
 - about 99
 - collaboration object messages 100
 - configuring 110
 - configuring connector agent logging destination 112
 - configuring InterChange Server Express logging destination 110
 - connector agent messages 102
 - connector messages 100
 - message formats 101
 - system logs 102
- logging on to System Monitor 13
- long-lived business processes 58
- lost connection to application 127

M

- managing
 - database connection pools 148
 - failures 123
 - log and trace files 103
 - Object Request Broker (ORB) 149
 - run-time properties 148
 - WebSphere MQ 147
 - WIP connections 123
- maps
 - administering 58
 - configuring tracing 115
 - operating 58
 - restart requirements 160
 - starting and stopping 60
 - states 59
 - tracing 106
 - viewing states in System Manager 59
 - viewing states in System Monitor 59
- meter display option, default monitors 13
- modifying schedules 82
- Monitor Definition Wizard, creating new monitors 6
- monitoring the system using InterChange Server Component Management view 25
- monitors
 - creating new 6
 - customizing the visual appearance in System Monitor 19

- monitors, default
 - bar display option 12
 - line display option 11
 - meterdisplay option 13
 - sample display options 10
 - stacked bar display option 12
 - table display option 10
 - table tree display option 11
 - using with System Monitor 2

O

- OAD, using as a Windows service 50
- Object Request Broker (ORB)
 - changing location 151
 - customizing 149
 - installing 149
 - managing 149
 - using the IBM Transient Name Server 152
- opening relationships 68
- operating
 - collaboration objects 52
 - connectors 42
 - InterChange Server Express 34
 - maps 58
 - relationships 60
- overriding schedules 81

P

- password
 - changing for database 36, 37
 - changing for InterChange Server Express 36
 - InterChange Server Express 36, 157
- persistent monitoring 23
 - accessing results 24
 - configuring 23
- persistent naming server 152

R

- refreshing System Manager 39
- relationship data
 - cleaning up participants 77
 - filtering displayed participants 76
 - printing relationship data 78
 - searching for participants 75
- Relationship Manager
 - starting 65
 - using 64
- relationships
 - administering 60
 - copying participants 73
 - creating relationship instances 71
 - deactivating and activating participants 72
 - deleting relationship instances 72
 - loading and unloading business object files 74
 - opening 68
 - operating 60
 - restart requirements 160
 - retrieving relationship instances 68
 - starting and stopping 62
 - states 61
 - table caching 63
 - working with 68
 - working with relationship data 74

- remote and automatic restart of connectors 48
- repos_copy
 - command 86
 - compiling components in the repository 94
 - deleting components from the repository 94
 - deploying a package to the repository 92
 - exporting components to a package 96
 - locale 97
 - printing a list of components in the repository 96
 - printing the command 92
 - syntax 87
 - usage scenarios 91
 - validating the package 92
 - validating the repository 94
- Repository (InterChange Server Express)
 - loading 86
 - repository, backing up 85
 - resolving
 - collaboration deadlocks 143
 - restart requirements 157
 - business objects 159
 - collaboration objects 158
 - collaboration templates 157
 - connectors 158
 - InterChange Server Express 157
 - maps 160
 - relationships 160
 - restarting connectors 48
 - run-time properties, configuring for collaboration objects 55
 - running, stopping, and pausing connectors 44

S

- sample display options, default monitors 10
- schedules
 - creating 81
 - deleting 83
 - displaying 83
 - enabling or disabling 83
 - modifying 82
 - overriding 81
- scheduling
 - collaboration objects 80
 - connectors 80
 - jobs 78
- server memory, controlling usage 121
- shutting down
 - connectors 47
 - InterChange Server Express 35
 - InterChange Server Express gracefully 35
 - InterChange Server Express immediately 36
 - System Manager 39
- stack trace, configuring 112
- stacked bar display option, default monitors 12
- starting
 - connectors manually 45
 - Flow Manager 136
 - InterChange Server Express 34
 - Relationship Manager 65
 - System Manager 39
- starting and stopping
 - maps 60
 - relationships 62
- starting, stopping, and pausing
 - collaboration objects 54
- startup parameters
 - customizing for InterChange Server Express 34

- startup parameters *(continued)*
 - InterChange Server Express 34
- state change log, deleting in System Monitor 22
- states
 - capturing state changes in System Monitor 22
 - collaboration objects 52
 - connectors 42
 - maps 59
 - relationships 61
- statistics
 - collaboration objects 28
 - connectors 30
 - InterChange Server Express 26
 - viewing and using from System Manager 26
 - viewing and using from the InterChange Server Component Management view 26
- System Manager
 - InterChange Server Component Management view 25
 - refreshing 39
 - restarting connectors 48
 - shutting down 39
 - shutting down connectors 48
 - starting 39
 - updating components 39
 - using 38
 - viewing collaboration object states in 53
 - viewing connector states in 43
 - viewing map states 59
 - viewing relationship states 61
- System Monitor 15
 - configuring existing views 18
 - creating new monitors 6
 - creating your own views 17
 - deleting historical statistics 22
 - deleting the business object probe data file 22
 - deleting the state change log 22
 - deleting views 18
 - fine-tuning 18
 - capturing state changes 22
 - customizing the visual appearance of monitors 19
 - resetting run-time statistic values 21
 - setting the default view 19
 - setting the frequency for capturing historical data 20
 - setting the refresh rate for run-time values 20
 - interface 15
 - logging on 13
 - restarting connectors 48
 - shutting down connectors 47
 - using default monitors 2
 - using display options 9
 - using views 16
 - viewing collaboration object states 52
 - viewing connector states 42
 - viewing map states 59
 - viewing relationship states 61

T

- table caching, relationships 63
- table display option, default monitors 10
- table tree display option, default monitors 11
- trace file management 103
- tracing 104
 - business objects (flow tracing) 106
 - collaboration objects 105
 - configuring 110
 - configuring collaboration object trace levels 114

- tracing *(continued)*
 - configuring connector tracing 114
 - configuring InterChange Server Express tracing destination 110
 - configuring map tracing 115
 - configuring trace levels 113
 - connectors 105
 - InterChange Server Express 105
 - maps 106
 - using 115
 - using log viewer 115
 - Web gateways 106
- transactional collaborations, failed 128
- troubleshooting 99
- typographic conventions vii

U

- unknown status, connector agent 127
- unresolved flows
 - finding 138
 - handling 143
 - viewing details 141
- updating components in System Manager 39

V

- viewing
 - connector states 42
- viewing log messages 115
- views
 - configuring existing in System Monitor 18
 - creating in System Monitor 17
 - deleting in System Monitor 18
 - setting the default in System Monitor 19
 - using default in System Monitor 16
 - using in System Monitor 16

W

- Web gateways, tracing 106
- WebSphere MQ
 - commands 155
 - errors 147
 - failure recovery and the queue manager 147
 - failures to create new instances of MQQueueManager 147
 - managing 147
 - memory checker segmentation fault 147
- WIP connections, managing 123
- working with relationships 68



Printed in USA