

IBM WebSphere InterChange Server



# User Guide

*V4.2.1*



IBM WebSphere InterChange Server



# User Guide

*V4.2.1*

**Note!**

Before using this information and the product it supports, read the information in "Notices" on page 83.

**30September2003**

This edition of this document applies to IBM WebSphere InterChange Server, version 4.2.1, WebSphere Business Integration Adapters, version 2.3.1, and to all subsequent releases and modifications until otherwise indicated in new editions.

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

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

© **Copyright International Business Machines Corporation 1997, 2003. All rights reserved.**

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

---

# Contents

<b>About this document</b> . . . . .	<b>v</b>
Audience . . . . .	v
Related documents . . . . .	v
Typographic conventions . . . . .	v
<b>New in this release</b> . . . . .	<b>vii</b>
New in IBM WebSphere Business Integration Express for Item Synchronization V4.3 . . . . .	vii
<b>Chapter 1. Overview</b> . . . . .	<b>1</b>
Architecture and components. . . . .	1
Toolset . . . . .	2
System Manager . . . . .	2
Summary of tasks . . . . .	3
<b>Chapter 2. Configuration</b> . . . . .	<b>5</b>
Using connector configurator . . . . .	5
Distributed connector agents . . . . .	5
Distributing a connector agent . . . . .	6
<b>Chapter 3. Startup</b> . . . . .	<b>9</b>
<b>Chapter 4. Monitoring the system</b> . . . . .	<b>11</b>
Using the Web-based System Monitor . . . . .	11
Using persistent monitoring . . . . .	31
<b>Chapter 5. Operating components of the system.</b> . . . . .	<b>35</b>
Overview of operating the system. . . . .	35
Operating InterChange Server Express . . . . .	35
System Manager tasks . . . . .	40
Operating connectors . . . . .	43
Operating collaboration objects . . . . .	53
Operating maps . . . . .	57
Operating relationships . . . . .	58
Backing up IBM WebSphere ICS components . . . . .	61
Using Repos_Copy . . . . .	64
Scheduling jobs . . . . .	76
<b>Notices</b> . . . . .	<b>83</b>
Programming interface information . . . . .	84
Trademarks and service marks . . . . .	84



---

## About this document

IBM<sup>(R)</sup> WebSphere<sup>(R)</sup> Business Integration Express for Item Synchronization V4.3 (also referred to in this document as Business Integration Express for Item Synchronization) provides a solution for enterprises that need to exchange business data with trading partners who require adherence to UCCnet standards.

This document provides an overview of the product and describes how to monitor and operate the system after you have installed it.

---

## Audience

This document is for system administrators, consultants, and developers who administer Business Integration Express for Item Synchronization.

---

## Related documents

The following additional documentation is available [actual document titles to come]:

- Installation Guide  
Describes how to install the product.
- Quick Start Guide  
Provides a brief overview of the tasks you will need to perform to get your system up and running.
- Guide to the UCCnet collaboration  
Describes the details [to come]
- Context-sensitive Help  
Context-sensitive help, activated when you press F1 within any tool, provides the detailed instructions for using that tool.
- Adapter guides  
An individual guide is provided for the use of each of the adapters utilized in a Business Integration Express for Item Synchronization solution.

You can install the documentation from the following site: [URL to come]

---

## Typographic conventions

This document uses the following conventions:

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

---

---

[ ]	In a syntax line, square brackets surround an optional parameter.
...	In a syntax line, ellipses indicate a repetition of the previous parameter. For example, <code>option[,...]</code> means that you can enter multiple, comma-separated options.
< >	In a naming convention, angle brackets surround individual elements of a name to distinguish them from each other, as in <code>&lt;server_name&gt;&lt;connector_name&gt;tmp.log</code> .
/, \	In this document, backslashes (\) are used as the convention for directory paths. For UNIX installations, substitute slashes (/) for backslashes. All product pathnames are relative to the directory where the Iproduct is installed on your system. Paragraphs inside a box with a UNIX or Windows label indicate notes listing operating system differences.

UNIX/Windows
--------------

<i>%text%</i> and <i>\$text</i>	Text within percent (%) signs indicates the value of the Windows text system variable or user variable. The equivalent notation in a UNIX environment is <i>\$text</i> , indicating the value of the <i>text</i> UNIX environment variable.
<i>ProductDir</i>	Represents the directory where the product is installed.

---



---

## **New in this release**

---

### **New in IBM WebSphere Business Integration Express for Item Synchronization V4.3**

This is the first release of this guide.



---

## Chapter 1. Overview

This chapter provides a high-level introduction to the use of IBM WebSphere Business Integration Express for Item Synchronization V4.3 (also referred to in this guide as Business Integration Express for Item Synch) as a comprehensive solution for exchanging data between UCCnet-compliant trading partners who are using different applications at their respective enterprise sites. Business Integration Express for Item Synchronization provides the capability to

- Process and route business information among disparate applications in the enterprise environment
- Exchange business information between trading partners who use disparate applications across the Internet.

This chapter provides an overview of the architecture and components of Business Integration Express for Item Synchronization, and provides a summary of the tasks necessary for implementing the system.

**Note:** Throughout this manual, illustrations are examples only, used to show structure and concepts. They do not necessarily document specific actual components.

---

### Architecture and components

The Business Integration Express for Item Sync system uses a central infrastructure (InterChange Server Express) and modular components in a hub-and-spoke design.

In this design, business-process logic resides in *collaborations* at the hub; data is exchanged between the hub and the spokes in the form of *business objects*. Connectors, supply connectivity to applications (or to web servers or other programmatic entities) at the spokes, and the *Access Interface* makes it possible for remote spoke sites that do not implement WebSphere InterChange Server to make calls over the Internet to a hub site that does.

The components of the system include:

- InterChange Server Express

An integration broker that functions as the hub for data exchanges

- Item Synchronization collaboration

**Collaborations** are software modules that contain logic that describes a distributed business process. Collaborations coordinate the functionality of business processes for disparate applications and enable data exchange between them. Collaborations are the hub; through them, data is exchanged in the form of business objects.

- Adapters

Each **adapter** includes a component that provides connectivity (referred to as a **connector** throughout this guide) and (for many adapters) a mechanism that generates or provides application-specific objects that are used by that connector. Some adapters, such as the Adapter for JText, are typically used for exchanging data with back-end applications at your enterprise site. Other adapters, such as the e-Mail Adapter, are typically used for data exchanges across the Internet with a UCCnet trading partner.

These connectors transform data from the application into business objects that can be manipulated by the collaborations, and transform business objects from the collaborations into data that can be received by the specific application. Each connector consists of two parts—the **connector controller** and the **connector agent**. The connector controller interacts directly with WebSphere InterChange Server collaboration objects and resides on a server that has implemented InterChange Server Express (the hub in a hub-and-spoke relationship). The connector agent interacts directly with an application, and can reside with that application on any server.

- Business Objects

Business objects are the messages used by the Business Integration Express for Item Sync system for exchanging data.

- Maps

Maps are used between a business object that is structured for the data model of a specific application and a business object that is generically structured for use by collaborations at the hub.

- Access Interface enables calls from external sites.

The Access Interface is part of the InterChange Server. It is a CORBA-compliant API that accepts synchronous data transfers from either internally networked or external sources. The data is then transformed into business objects that can be manipulated by a collaboration. The Access Interface makes it possible to receive calls from external entities—for example, from web browsers at remote customer sites—that do not come through connector agents, but instead come through web servlets into the Server Access Interface [OK?].

---

## Toolset

Available with the business integration system are development tools, administrative tools, and System Manager—which helps you generally with using your system and accessing the tools, and which provides an interface for deploying the integration projects that you develop.

---

## System Manager

The System Manager interface is derived from the Eclipse platform—an open-source integrated development environment for the creation of tools. The Eclipse platform provides tools developers with a development kit and runtime that enables the developer to write plug-ins that allow the user to work with a particular type of resource.

System Manager is installed as a plug-in that can be used in IBM branded versions of the Eclipse platform.

You access System Manager as an Eclipse-based *perspective*. The System Manager perspective provides views and editors for

- Accessing tools
- Performing configuration tasks on individual modular components
- Working with instances of InterChange Server
- Handling groups of integration components as user projects, and deploying components to the repository

Tasks in System Manager are accessible through a menu and corresponding toolbar, as well as through the System Manager view. The System Manager view,

displayed in the left-hand panel of System Manager, shows a folder-style list of libraries of modular integration components (such as business object definitions, database connection pools, and collaboration templates and objects). You can perform development and configuration tasks by right-clicking on the icon for a component in a library.

## Tools

You can use the provided tools to configure connectors, and create and modify business objects, maps, and relationships. [OK?] lists and describes some of the software tools provided.

Table 1. Tools

Tool	Description
Test Environment and Test Connector	Provides an environment in which you can test business integration interfaces you have developed. Provides graphical interfaces to emulate connectors, start the required components, and examine business object data.
Connector Configurator	Used for adding application-specific properties to a connector definition, for setting property values, and for configuring the connector definition with its business objects and maps
Map Designer	A graphical tool that specifies data transformations between application-specific and generic business objects.
Relationship Designer	A graphical tool that defines relationships between types of objects. These relationships are important in mapping, for example, to specify the relationship between one type of business object and another.
Business Object Designer	A forms-based interface used for creating business object definitions both manually and from Object Discovery Agents (ODAs)

## Summary of tasks

To implement Business Integration Express for Item Synch, you will typically need to perform the following tasks:

1. Set up the database you will use  
You can use either DB2 [version to come] or MS SQL [version to come]. If you do not already have either database software, the Installer can optionally install a version of DB2 for you.
2. Run the Installer  
The Installer automates many tasks that would otherwise require manual configuration. When you run the Installer, you make certain selections, including designating which database software you will use and which adapters you will use. The Installer then does much of the installation and configuration for you. For details on running the Installer and the selections you can make, see the *Installation Guide*.
3. Choose your adapters  
Make your adapter selections when you run the Installer. If you later decide to add another adapter, you can run the Installer again, and select the additional adapter.
4. Create application-specific business objects  
Generic business objects are provided with the Item Synchronization Collaboration. To implement your solution, you will need to create application-specific objects for your connector. You can do so by using Business Object Designer, or by using the Object Discovery Agent (ODA) if one is provided with your adapter. For details about this task, start Business Object Designer, and press F1 to display the context-sensitive help.

5. Create maps between the application-specific business objects and the generic business objects used by the collaboration

The process for creating maps is described in the context-sensitive help for the Map Designer.

6. Configure your adapters

To configure your adapters, you will need to run Connector Configurator and set certain connector properties, and designate the business objects that the adapter will support. These tasks are described in detail in the Connector Configurator appendix included with each individual adapter guide. You may also need to perform certain set up tasks that are specific to the application to which the adapter connects, as described in guide for that specific adapter.

You may need to do additional configuration if you wish to use implementations of the adapter other than the default. For example, by default, the Installer sets up the adapter's connector agent to run locally, on the same machine on which InterChange Server Express is installed. [TRUE?] If you wish to run the connector agent on a different machine, you must do additional configuration, as described in "Distributed Connector Agents," later in this guide.

7. [IS THERE ANOTHER TASK HERE, FOR "CONFIGURING THE HOST" --OR IS THAT SPECIFIC TO IMPLEMENTING A DISTRIBUTED CONNECTOR AGENT ON A DIFFERENT MACHINE THAN THE ICS?]

8. Start the system

For this task, see "Startup," later in this guide.

---

## Chapter 2. Configuration

The Installer provided with the product performs most configuration tasks for you. However, you may still need to perform certain tasks to configure your adapters.

You must set properties for your connector, and designate the business objects that it supports. [TRUE, OR DOES THE INSTALLER TAKE CARE OF THIS?] These tasks are described in "Using connector configurator." You may also need to perform certain set up tasks that are specific to the application to which the adapter connects. Those tasks, if necessary, are described in the guide for the specific adapter.

You may need to do additional configuration if you wish to use implementations of the adapter other than the default. For example, by default, the Installer sets up the adapter's connector agent to run locally, on the same machine on which InterChange Server Express is installed. [TRUE?] If you wish to run the connector agent on a different machine, you must do additional configuration, as described in Distributed Connector Agents.

---

### Using connector configurator

[THE CONTENT OF THIS SECTION WILL BE A REVISED VERSION OF THE CONNECTOR CONFIGURATOR APPENDIX THAT CURRENTLY RESIDES IN EACH ADAPTER GUIDE. TO COME.]

---

### Distributed connector agents

[I'M UNCERTAIN OF THE CONTENT WE SHOULD USE TO DESCRIBE THE SCENARIO WHERE A CONNECTOR AGENT IS ON A DIFFERENT MACHINE THAN THE ICS, BUT IS NOT ACROSS AN INTERNET FIREWALL (AS WOULD BE THE CASE FOR A REMOTE AGENT). THE FOLLOWING TEXT IS FROM THE 4.2.1 SYSTEM IMPLEMENTATION. IS IT VALID AND USEFUL FOR ITEM SYNC?]

Although it is possible to install and run IBM WebSphere Business Integration Adapters on the same computer where InterChange Server is installed and running, it is sometimes necessary to distribute connector agents. This means installing the connector agent on a different computer than the computer on which ICS is installed.

There are several reasons you might have to distribute connectors agents, including the following:

- Connector performance often improves if the connector agent is installed on the computer that hosts the application with which it communicates.
- Some connectors must be installed on the computer where the application they communicate with is installed.
- A connector must be installed completely outside of the network in which InterChange Server runs. For example, if a site has an arrangement with a trading partner that allows the partner's applications to integrate with the site's applications then you must install the agent within the partner's network and enable it to communicate with the site through Remote Agent technology.[IS THIS CORRECT FOR ITEM SYNC?]

---

## Distributing a connector agent

The following sections describe the tasks necessary to distribute a connector agent:

- “Prepare the agent host computer”
- “Copy required files and libraries”
- “Modify the connector agent environment”

### Prepare the agent host computer

Perform the following prerequisite steps to prepare the agent host computer for the tasks that follow in subsequent sections:

1. Create a directory for the Business Integration Express for Item Sync software on the agent host computer.
2. Create a subdirectory named `connectors` within the directory created in step 1.
3. Create a subdirectory named `messages` in the `connectors` directory created in step 2.

### Copy required files and libraries

Connectors depend on several files and directories to execute properly. Perform the following steps to properly migrate a connector agent:

1. Copy the following files or directories from the IBM WebSphere Business Integration Adapter directory on the InterChange Server [add Express] host computer to the corresponding directory on the agent host computer:
  - `InterchangeSystem.cfg` configuration file
  - `messages` directory
  - `bin` directory
  - `lib` directory
2. Migrate the connector agent libraries by copying the entire directory for the connector agent from the `connectors` directory on the InterChange Server host computer to the corresponding directory on the agent host computer.

For example, if the connector agent being distributed was that for the WebSphere Business Integration Adapter for JDBC, then copy the `Jdbc` directory into the `\connectors` directory.
3. Copy the message file for the connector agent into the `\connectors\messages` directory within the product directory on the agent host computer.

Connector message files have the same name as the connector and have a `.txt` extension. For example, if the connector agent being distributed was that for the WebSphere Business Integration Adapter for JDBC, then copy the `JDBCConnector.txt` file into the `\connectors\messages` directory.

### Modify the connector agent environment

Connector agents require a number of pieces of information to run correctly, such as the name of ICS, and the location of the supporting libraries for WebSphere InterChange Server.

#### Modifying the location of the WebSphere InterChange Server directory

Connector agents need to know the location of the product directory to reference things such as the supporting classes that are required. Take one of the following two approaches to configure the startup mechanism of the connector agent so that it can reference the location of the WebSphere InterChange Server directory:



- Create a system-level environment variable named CROSSWORLDS that identifies the location of the WebSphere InterChange Server directory on the agent host computer. Connector agent startup scripts are typically designed to query for a system-level environment variable by that name, so you should not have to do further work to configure them to reference it.

The advantages of this approach are that if one computer was going to host multiple distributed connector agents, then all of their startup scripts could then reference the same environment variable, and that you do not have to modify the connector startup script for this particular purpose.

The disadvantage of this approach is that you must create the environment variable, which is an extra task.

- Edit the startup script for the connector agent and define a variable named CROSSWORLDS within the file that points to the WebSphere InterChange Server directory on the agent host computer.

The advantage of this approach is that the extra work of creating the environment variable is not required. Furthermore, if multiple instances of a single connector must be installed on multiple distributed computers then you only have to edit the file once and can use it on all of the necessary computers provided that the WebSphere InterChange Server directory is the same on all of them.

The disadvantage of this approach is that you must manually edit the connector startup script to create the variable within it.

Many of the other variables defined in the startup scripts make reference to the CROSSWORLDS variable (to ultimately reference dependency classes in the \lib directory, for instance), so the variable must be defined prior to any others that reference it. The lines below are from the start\_connector.bat file in a Windows environment; the second of the lines was added to define the CROSSWORLDS variable:

```
setlocal
set CROSSWORLDS=C:\CrossWorlds
REM Define local batch PATH to insure we execute our jre
set PATH="%CROSSWORLDS%\bin;%PATH%
```

## Name of connector agent and InterChange Server instance

For a connector agent to communicate with ICS, you must configure it to reference ICS by the proper case-sensitive name. Similarly, you must configure the agent so that when it starts, it is assigned the proper case-sensitive name for the connector as defined in the ICS repository. Take one of the following three approaches to configure the startup mechanism of the connector agent so that it can reference the location of the WebSphere InterChange Server directory:

- Start the connector by executing the appropriate startup script and manually typing the name of the connector agent and the name of ICS each time.

The advantage of this approach is that you can typically use a single connector startup script for multiple connectors, because you manually enter the variable information each time you execute the script.

The disadvantage of this approach is that the responsibility of passing the variables with the proper case and spelling is on you.

This approach is typically used in a Solaris environment—you execute the connector\_manager.sh script and pass the name of the connector agent and ICS at the command line while doing so. This is also a viable approach in a Windows environment, though less common because shortcuts make it easier to pass the variable information with less of a threat of misspelling the case-sensitive names.

- Start the connector by hard-coding the name of the connector agent and the name of ICS into the startup script and executing it each time. Below is an example where the start\_JText.bat file on a Windows computer is modified to enable startup of the JTextConnector by hard-coding the variable information; note that the CONNDIR, CONNAME, and SERVER variables all required modification:

```
setlocal
set CROSSWORLDS=C:\CrossWorlds
REM Define local batch PATH to insure we execute our jre
set PATH="%CROSSWORLDS%\bin;%PATH%"
REM set the directory where the specific connector resides
set CONNDIR="%CROSSWORLDS%\connectors\JText"
REM goto the connector specific drive & directory
cd /d %CONNDIR%
REM set the name to be the application connector that is starting
set CONNAME=JText
REM set the server name to be the interchange that is being targeted
set SERVER=CrossWorlds
```

The advantage of this approach is that you can ensure the case and spelling of the variables and minimize the amount of manual effort involved in starting the connector agent.

The disadvantage of this approach is that it is very customized to a particular agent; if you need to start multiple agents, then you must customize a startup script for each one.

- Include the name of the connector agent and the name of ICS in the target of a shortcut you execute, which ultimately executes the appropriate startup script. The advantage of this approach is that you can ensure the case and spelling of the variables and minimize the amount of manual effort involved in starting the connector agent, just as when hard-coding the variables within the script, but without creating a custom script for each agent. The disadvantage of this approach is that it is not viable in all environments: for instance, in a situation where the connector agent must operate as a service a shortcut cannot be used; some connector agents require their own custom startup scripts anyway (typically Java-based connectors), so there is potentially no benefit in that respect.

## Installing the IBM WebSphere MQ client on the agent host computer

Connector agents require the WebSphere MQ client. The client must be installed on the agent host computer.

---

## Chapter 3. Startup

To start your system after it has been installed and configured, follow this sequence:

1. Start the Name Server.
2. Start InterChange Server Express.
3. Start the adapters. [IS THIS A REQUIRED STEP? WHEN YOU START ICS, DOES THAT ALSO START THE CONNECTOR CONTROLLER? DOES IT START THE CONNECTOR AGENT? IS IT DIFFERENT IF THE CONNECTOR AGENT IS ON A DIFFERENT MACHINE THAN THE ICS?]

To start Name Server and InterChange Server Express, use the provided \*.bat files. You can start both using a single .bat file, or start each separately:

- To start both Name Server and InterChange Server Express in the correct order, use start\_server\_service.bat. [CORRECT?]
- To start only Name Server, use [WHAT NAME?]
- To start only InterChange, use start\_server.bat [CORRECT?]

For details about starting the adapters, see "Operating Connectors," later in this guide. For shutting down and other operational details about InterChange Server Express, see "Overview of operating the system," later in this guide.



---

## Chapter 4. Monitoring the system

[EVERYTHING IN THIS CHAPTER IS FROM THE 4.2.1 SYSTEM ADMINISTRATION GUIDE. IS IT ALL VALID FOR ITEM SYNC? THE FOLLOWING TOPICS HAVE BEEN REMOVED FROM THE ORIGINAL CHAPTER:

- SNMP
- UNIX SPECIFIC INFORMATION
- WINDOWS-BASED SYSTEM MONITOR]

Monitoring includes monitoring all IBM WebSphere ICS components, such as connectors and collaboration objects, as well as the connection to all integrated applications. The primary tool for this is Web-based System Monitor.

---

### Using the Web-based System Monitor

The Web-based System Monitor is a tool that allows you to monitor the Business Integration Express for Item Sync system from the Web. It allows you to configure how you view the data and also allows you to view historical data in addition to current data. It allows you to start, stop, and pause components.

#### Different approaches for using the Web-based System Monitor

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

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

**Monitors** are definitions of the information you want to view. You can either create new monitors using the Monitor Definition Wizard, a tool launched from System Manager, or you can use the default monitors provided with the Monitor Definition Wizard. See Table 2 on page 12 for a description of the default monitors provided with the Monitor Definition Wizard. **Views** are definitions of the monitor or monitors you want displayed when you monitor the system using the Web-based System Monitor. You can create and configure views from the Web-based System Monitor, or you can use the default views provided with the Web-based System Monitor. See Table 4 on page 25 for a description of the default views included with the Web-based System Manager. Once you have created the monitors and views you need, or decided to use the default monitors and views, you can then monitor the ICS system from the Web.

#### Using default monitors

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

Table 2 on page 12 contains a description of the default monitors provided with the Monitor Definition Wizard.

Table 3 on page 15 contains a description of display options listed in Table 2.

“Sample display options” on page 15 contains samples of the display options listed in Table 2.

*Table 2. Default monitors*

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

Table 2. Default monitors (continued)

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

Table 2. Default monitors (continued)

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



Table 2. Default monitors (continued)

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

Table 3 describes the configurable properties of each display option listed in Table 2 on page 12. To see examples of the display options, see “Sample display options.”

Table 3. Configurable properties of display options

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

### Sample display options

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

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

## Connector Statistics



Display  items

1-9 of 9 items

Select all displayed items

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

Sat Mar 15 16:16:21 PST 2003

Figure 1. Table display option

## System Overview



Display  items

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

Select all displayed items

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

Figure 2. Table tree display option




Figure 3. Line display option




Figure 4. Bar display option




Figure 5. Stacked bar display option

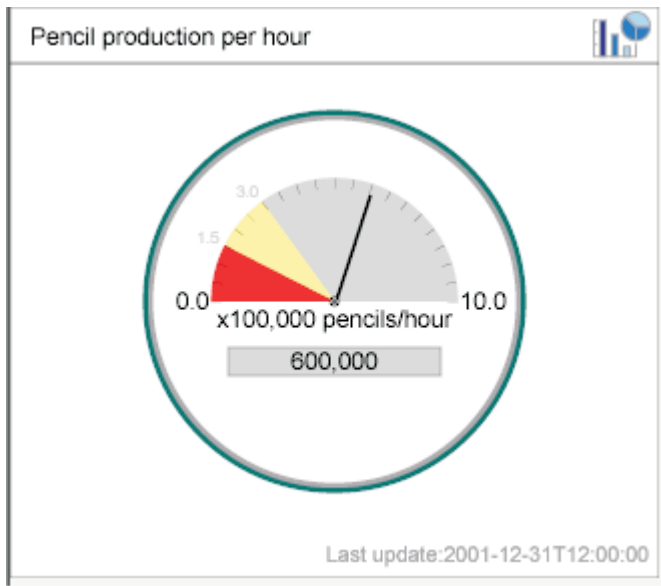


Figure 6. Meter display option

## Creating new monitors using the Monitor Definition Wizard

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

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

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

**Note:** Before creating new monitors, be sure to look at the existing default monitors in Table 2 on page 12, to see if the monitor you want to create already exists.

[STILL OK?] To create a new monitor, follow these steps:

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

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

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

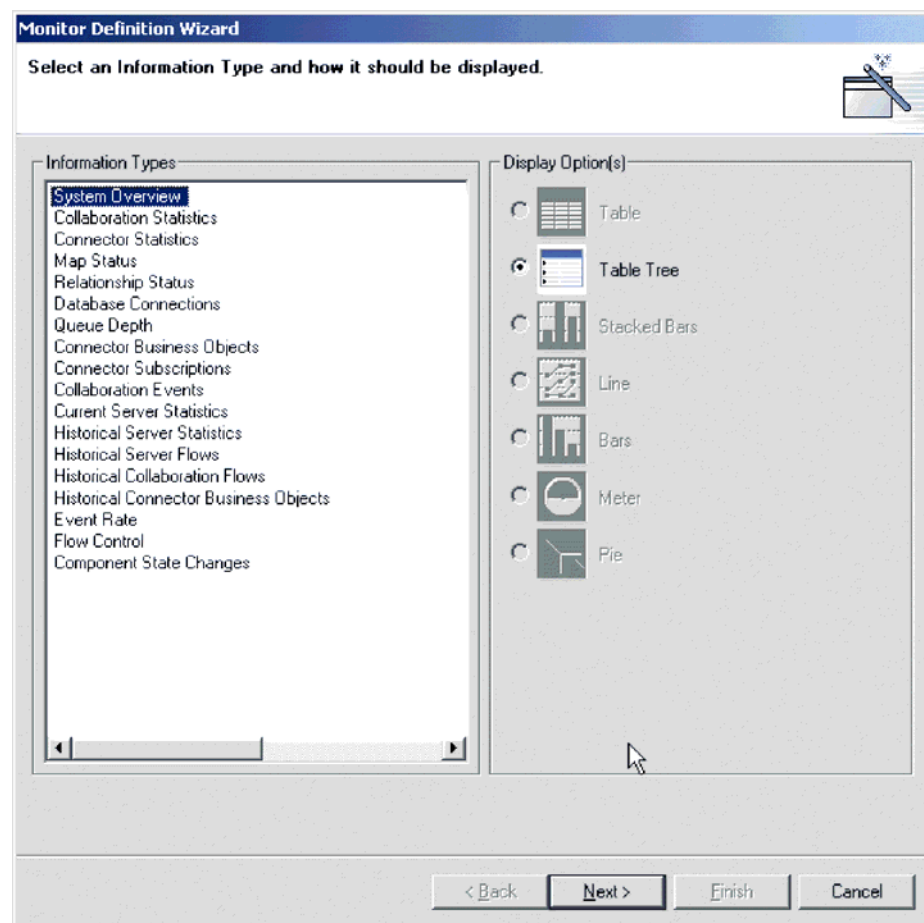


Figure 7. Monitor Definition Wizard, screen for selecting information type and display option

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

available display options. To see examples of the display options, see “Sample display options” on page 15, and for a description of the configurable properties of each display option, see Table 3 on page 15..

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

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

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

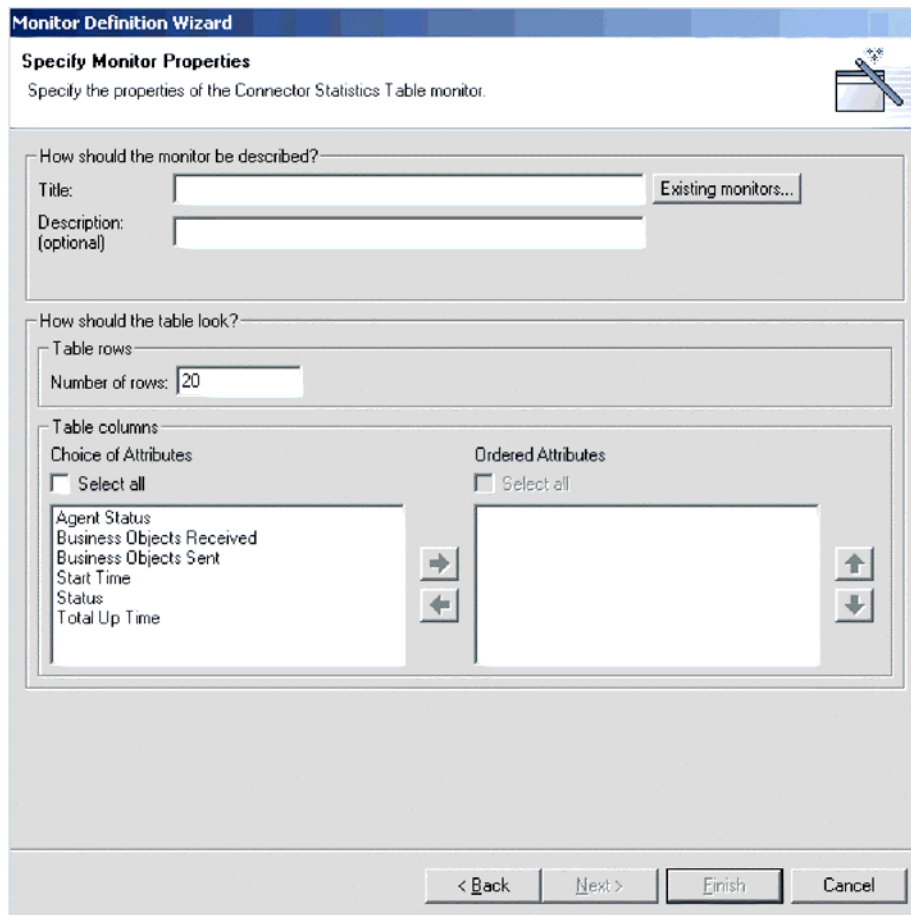


Figure 8. Monitor Definition Wizard, Specify Monitor Properties screen

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

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

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

For a list of all display option configurable properties, when using either the Monitor Definition Wizard or the Web-based System Monitor, see Table 3 on page 15..

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

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

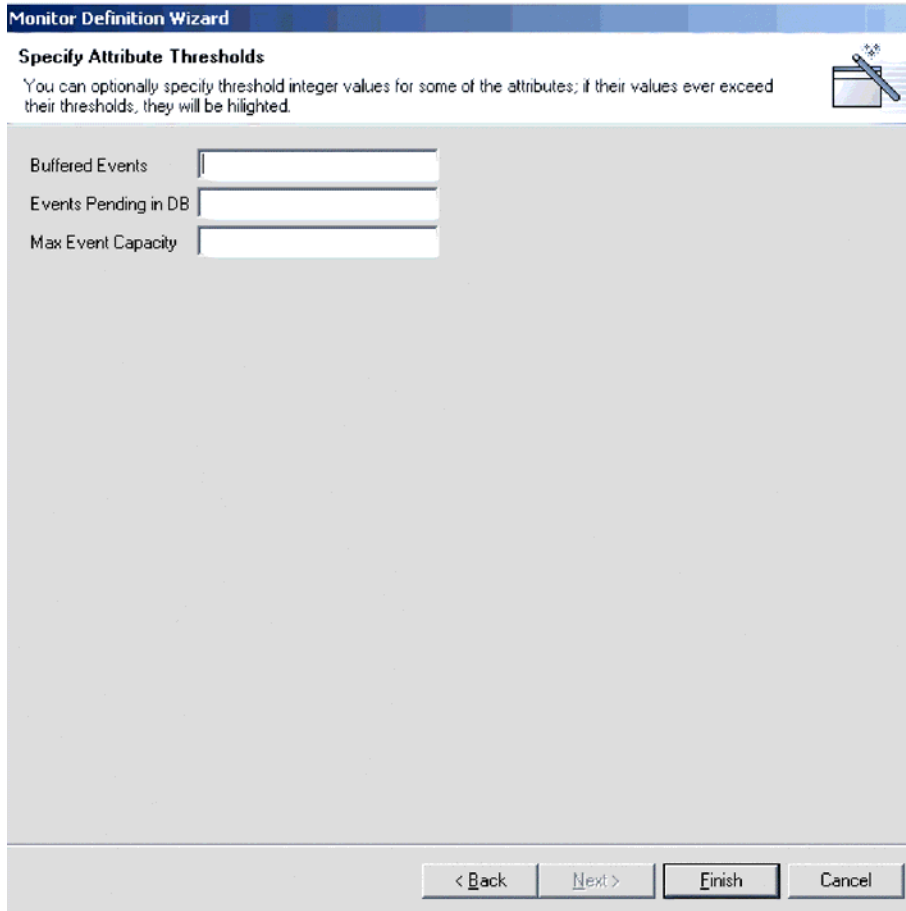


Figure 9. Monitor Definition Wizard, Specify Attribute Thresholds screen

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

## Logging on to the Web-based System Monitor

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

- Start InterChange Server (ICS) on the machine being monitored.
- Make sure the Web-based System Monitor and the application server are installed.
- Start the application server.
- Obtain the username and password necessary for logging on to the Web-based System Monitor. The username and password are the same as those used when logging on to ICS.

To log on to the Web-based System Monitor, do the following:

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



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

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

`http://monitorserver/ICSMonitor`

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

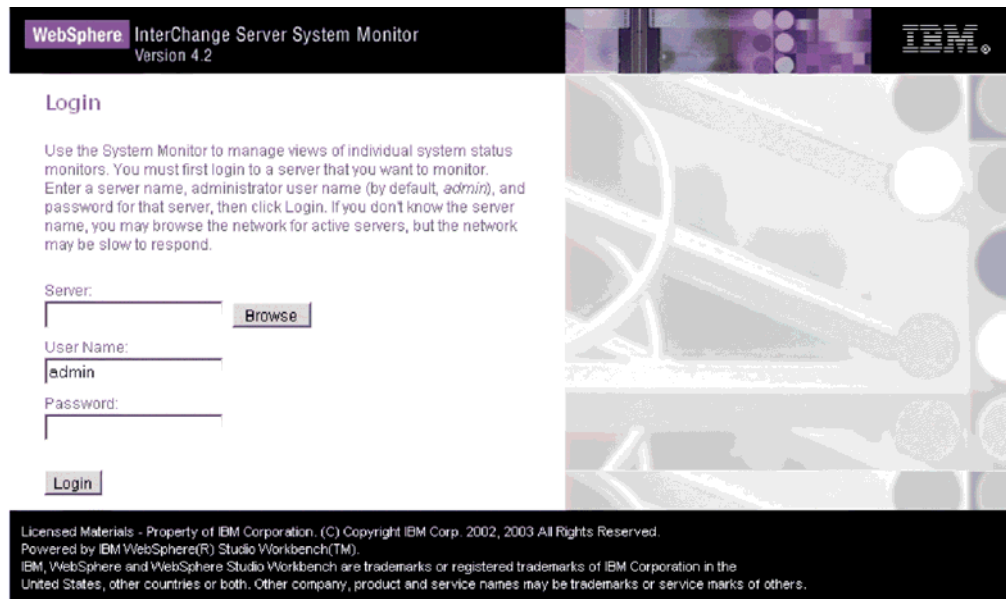


Figure 10. Web-based System Monitor, Login screen

2. Type the server name, user name, and password for the InterChange Server instance you want to monitor, then click Login. The Web-based System Monitor appears (see Figure 11 on page 24).

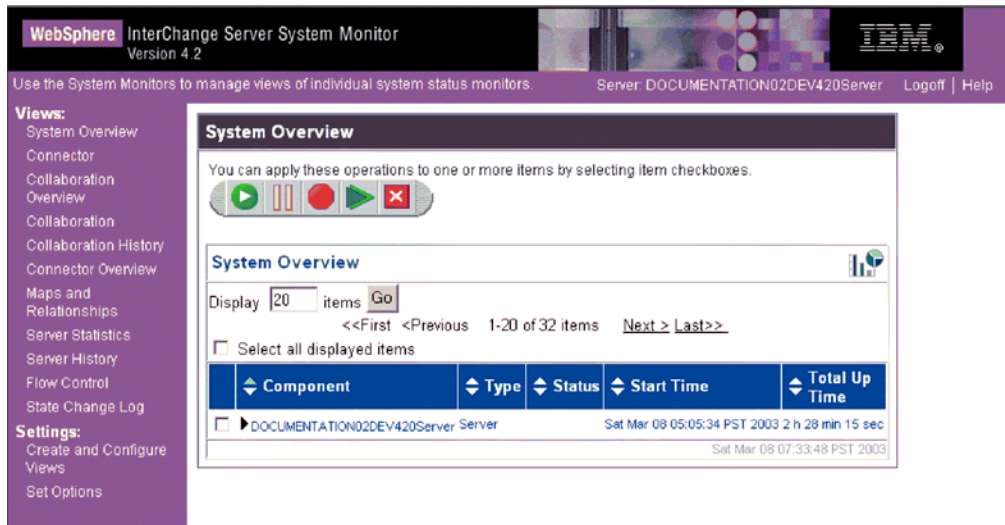


Figure 11. Web-based System Monitor, displaying System Overview as the default view

## Contents of the Web-based System Monitor Web page

The Web-based System Monitor Web page contains the following items:

- **List of views:** Initially, the views listed in the left column under Views are the default views provided with the installation of the Web-based 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 26), 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 the Web-based System Monitor. For instructions on creating, configuring, or deleting views, or setting the default view you see when you log on, refer to the following sections:
  - “Creating new views” on page 26
  - “Configuring existing views” on page 27
  - “Deleting views” on page 27
  - “Setting the default view” on page 28
- **Set Options link:** The Set Options link opens the Set Options dialog box (see Figure 14 on page 29), which allows you to set the following system-wide or component settings:
  - Refresh rate of the views that display current statistics
  - Frequency of captured historical data captured for each component type
  - Reset component statistics to “0”
  - Capture component state changes
  - Delete component state change log
  - Delete historical statistics for all components in the system
  - Delete business object probe data log

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

- “Setting the refresh rate for runtime values being monitored” on page 29
- “Setting the frequency for capturing historical data” on page 29

- “Resetting runtime statistic values” on page 30
- “Capturing state changes” on page 31
- “Deleting the state change log” on page 31
- “Deleting historical statistics” on page 31
- “Deleting the business object probe data file” on page 31
- **Default view:** A default view is displayed in the body of the Web page when you log on to the Web-based System Monitor. The first time you open the tool, the System Overview view included with the product is displayed. To change the default view displayed, see “Setting the default view” on page 28.
- **Logoff link:** The Logoff link allows you to log off the Web-based System Manager.
- **Help link:** The Help link opens an HTML page with the following information:
  - A link to download the documentation set for the Business Integration Express for Item Sync product
  - A directory location on your local machine where you can launch the Table of Contents file with links to help topics. This assumes you have already downloaded the documentation set.

## Using Views

You can either begin monitoring the system using the default views, or you can add, change, or delete views before monitoring the system. The following sections describe how to use existing views or create and configure views from the Web-based System Monitor. A **view** is a Web page that contains one or more monitors. Several default views are included in the installation of the Web-based System Monitor. You may use these default views or create new views. Before you can create and configure views, you must log on to the Web-based System Monitor. For instructions on logging on to the Web-based System Monitor, see “Logging on to the Web-based System Monitor” on page 22.

This section covers the following topics:

“Using default views”

“Creating new views” on page 26

“Configuring existing views” on page 27

“Configuring existing views” on page 27

### Using default views

The default views that are shipped with the Web-based System Monitor are described in Table 4. The table describes which monitor or monitors are contained in the view, as well as which display option is used. For descriptions of default monitors and samples of display options, see Table 2 on page 12 and “Sample display options” on page 15, respectively.

*Table 4. Default views*

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

Table 4. Default views (continued)

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

## Creating new views

The following instructions describe how to create a new view:

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

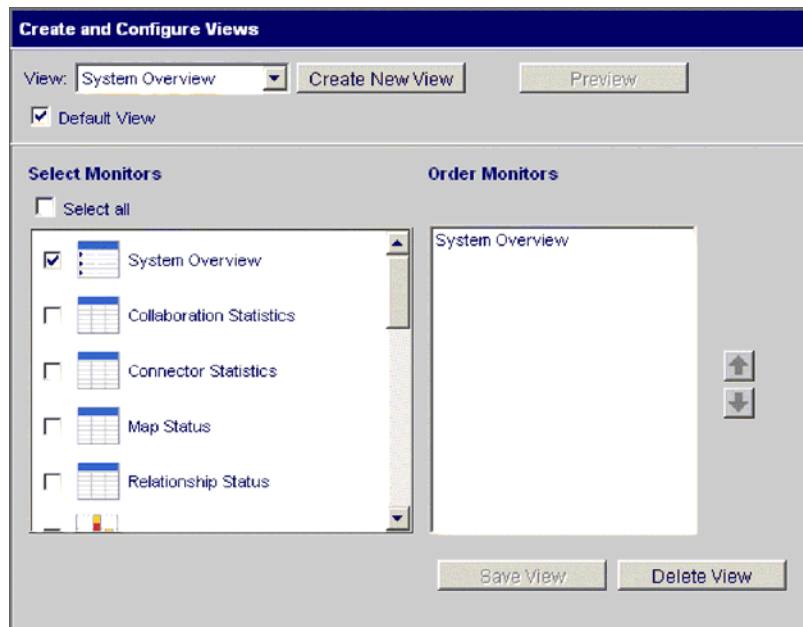


Figure 12. Web-based System Monitor, Create and Configure Views dialog box

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

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

### **Configuring existing views**

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

1. Click Create and Configure Views from the left frame of the Web-based System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 26).
2. Select the view you want to change from the View drop-down menu.
3. Add monitors to or remove monitors from the view, from the Select Monitors group box. The revised monitors for the view appear in the Order Monitors group box.
4. Use the up and down arrows to the right of the Order Monitors group box to put the monitors in the order you want to view them, from top to bottom.
5. Click Preview if you want to see a preview of the new view.
6. Click Save View. A "View was saved successfully" message appears.

### **Deleting views**

The following instructions describe how to delete a view:

1. Click Create and Configure Views from the left frame of the Web-based System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 26).
2. Select the view you want to delete from the View drop-down menu.
3. Click Delete View. A message appears, asking if you are sure you want to delete the view.
4. Click OK. The view is removed from the Views list in the left frame of the Web-based System Monitor.

## **Fine-tuning the Web-based System Monitor**

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

"Setting the default view" on page 28

"Customizing the visual appearance of the monitors" on page 28

"Setting the refresh rate for runtime values being monitored" on page 29

"Setting the frequency for capturing historical data" on page 29

"Resetting runtime statistic values" on page 30

“Capturing state changes” on page 31

“Deleting the state change log” on page 31

“Deleting historical statistics” on page 31

“Deleting the business object probe data file” on page 31

### Setting the default view

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

1. Click Create and Configure Views from the left frame of the Web-based System Monitor. The Create and Configure Views dialog box appears (see Figure 12 on page 26).
2. Select the view you want to be the default view from the View drop-down menu.
3. Select the Default View check box.
4. Click Save View. A “View was saved successfully” message appears. The next time you log on to the Web-based System Manager, the view you selected as the default view will be displayed.

### Customizing the visual appearance of the monitors

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

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

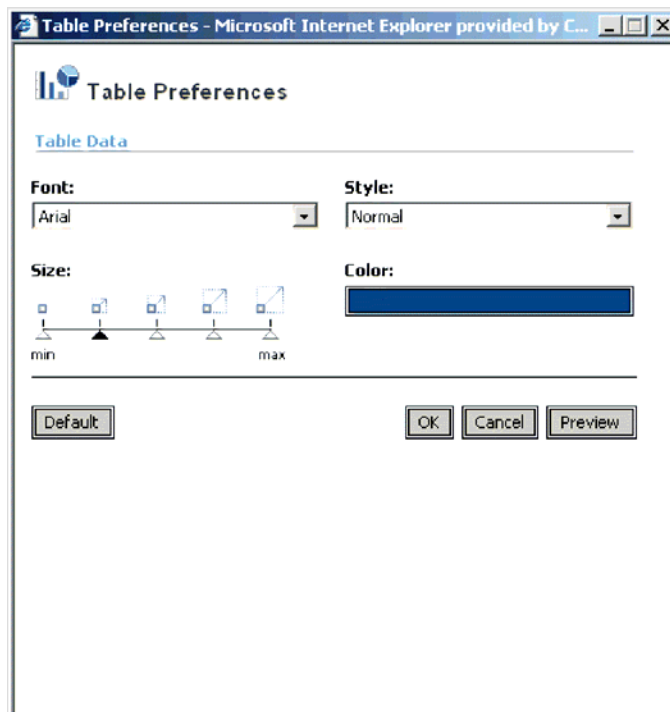


Figure 13. Web-based System Monitor, Table Preferences dialog box

- From the Preferences dialog box, select the appearance options you want to change. For a list of what appearance options are available with each display option, see the Runtime configurable properties column of Table 3 on page 15.
- Click Preview to see a preview of the changes you made.
- Click OK. The changes you made appear in the monitor.

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

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

### Setting the refresh rate for runtime values being monitored

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

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

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

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

Figure 14. Web-based System Monitor, Set Options dialog box

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

### Setting the frequency for capturing historical data

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

- Click Set Options from the left frame of the Web-based System Monitor. The Set Options dialog box appears (see Figure 14).
- 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 30).



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

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

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

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

### Resetting runtime statistic values

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

1. Click Set Options from the left frame of the Web-based System Monitor. The Set Options dialog box appears (see Figure 14 on page 29).
2. In the "Do you want to reset component statistics?" section, select the component from the Component Type drop-down menu.
  - If you select Server, then runtime statistics for all components are reset.



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

### **Capturing state changes**

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

1. Click Set Options from the left frame of the Web-based System Monitor. The Set Options dialog box appears (see Figure 14 on page 29).
2. Under the "Do you want to capture state changes of a particular component?" section, select the component from the Component Type drop-down menu.

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

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

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

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

1. Click Set Options from the left frame of the Web-based System Monitor. The Set Options dialog box appears (see Figure 14 on page 29).
2. Under the "Do you want to delete the state change log for all components?" section, click the calendar icons to enter the start date and end date for the data to be deleted, then click the Delete button in that section.

### **Deleting historical statistics**

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

1. Click Set Options from the left frame of the Web-based System Monitor. The Set Options dialog box appears (see Figure 14 on page 29).
2. Under the "Do you want to delete the historical statistics for all components?" section, click the calendar icons to enter the start date and end date for the data to be deleted, then click the Delete button in that section.

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

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

1. Click Set Options from the left frame of the Web-based System Monitor. The Set Options dialog box appears (see Figure 14 on page 29).
2. Under the "Do you want to delete the data for a business object probe?" section, select the business object probe from the Business Object Probe drop-down list, then click the calendar icons to enter the start date and end date for the data to be deleted, then click Delete.

---

## **Using persistent monitoring**

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

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

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

This section covers the following topics:

“Configuring persistent monitoring”

“Accessing the results from persistent monitoring” on page 33

## Configuring persistent monitoring

You configure the various levels of persistent monitoring for system components from the Web-based System Monitor.

[CORRECT, OR IS THIS WINDOWS BASED SYSTEM MONITOR ONLY?] You can also configure certain system-wide parameters of persistent monitoring from System Monitor. The following instructions explain how to do this.

1. Open the Edit Configuration tool by right-clicking the server from the Server Instances section of System Manager, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the InterchangeServer.cfg file.
2. Select the Misc tab (see Figure 16).

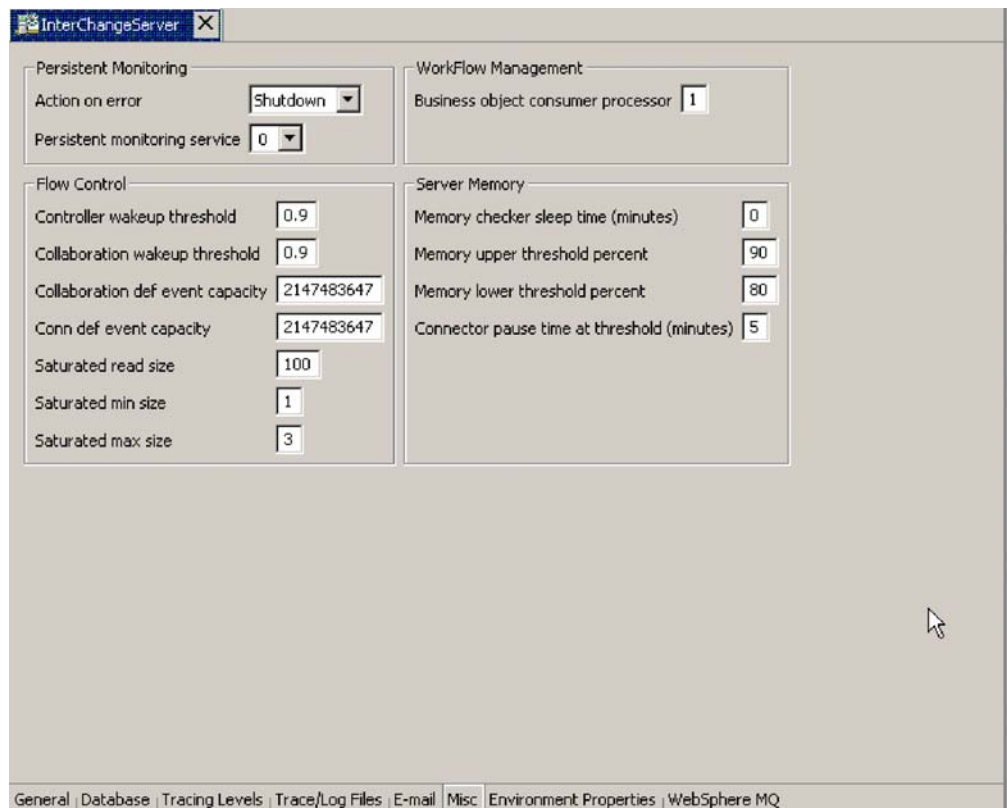


Figure 16. Edit Configuration screen, Misc tab

3. Do the following in the "Persistent Monitoring pane:

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

## **Accessing the results from persistent monitoring**

You access the results of persistent monitoring by using the Web-based System Monitor. Many default views are provided with the Web-based System Monitor that display historical state and statistical information. These include: Connector, Connector Overview, Collaboration Overview, Collaboration History, Server Statistics, and Server History. For more information on using default views, see "Using default views" on page 25. Alternatively, you can create your own views that can contain historical data. For more information on creating views, see "Creating new views" on page 26.



---

## Chapter 5. Operating components of the system

[EVERYTHING IN THIS CHAPTER IS FROM THE 4.2.1 SYSTEM ADMINISTRATION GUIDE. IS IT ALL VALID FOR ITEM SYNC? THE FOLLOWING TOPICS HAVE BEEN REMOVED FROM THE ORIGINAL CHAPTER:

- SNMP
- UNIX SPECIFIC INFORMATION
- WINDOWS-BASED SYSTEM MONITOR
- HA]

This chapter describes some of the tasks you may need to perform while operating an IBM WebSphere ICS system.

---

### Overview of operating the system

[OK?]When starting the Business Integration Express for Item Sync system, the order you must follow is to start WebSphere MQ Listener, Naming Server, InterChange Server, then IBM WebSphere System Manager. The connectors are automatically started when you start InterChange Server.

When shutting down an instance of InterChange Server 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 system, make sure all of the necessary third-party software is running. This includes the database on which the repository resides. This section assumes that you have already started the system and loaded the repository. [LOADING THE REPOSITORY NECESSARY THE FIRST TIME, OR IS IT DONE BY INSTALLER?]

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

---

### Operating InterChange Server Express

Operating InterChange Server Express may involve the following tasks:

“Starting InterChange Server Express” on page 35

“Shutting down InterChange Server Express” on page 37

“Changing the InterChange Server Express and database passwords” on page 38

### Starting InterChange Server Express

To start InterChange Server:

### UNIX

Run the `ics_manager -start` script.

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

### Windows

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

At startup, InterChange Server Express reads the `InterchangeSystem.cfg` file and sets its properties according to the parameter values listed there.

## Customizing InterChange Server Express startup parameters

To customize the InterChange Server Express startup:

### UNIX

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

`-start`

`-stop`

`-status`

### Windows

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

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

Table 5. InterChange Server Express startup parameters

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

## Shutting down InterChange Server Express

Shutting down InterChange Server Express stops all running collaborations and connectors, as well as InterChange Server Express itself. All connections to the database are closed and the machine's system resources used by InterChange Server Express are returned. Depending on your operating system, you can shut down the server using one of the following methods:

### UNIX

Run the `ics_manager -stop` script.

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

### Windows

In System Manager, right-click the server name under Server Instances in the InterChange Servers section, then select Shut Down > Gracefully.

Alternatively, you can select Shut Down > Immediately, which shuts down the server without cleanup. To determine which type of shutdown is best for you, refer to the following topics: "Graceful shutdown" on page 37 and "Immediate shutdown."

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

### Graceful shutdown

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

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

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

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

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

- InterChange Server Express shuts down.

**Note:** This procedure cannot be used to shut down an HA system. For information on shutting down a high-availability (HA) system, see on page 0. "Changing the status of a resource" on page 153

### Immediate shutdown

Immediately stopping the system forces the system to shut down without processing any more flows. Running connectors and collaborations are stopped immediately. When the system is restarted, flows that were interrupted by the immediate shutdown are redelivered in the same processing order. If one of these flows wrote data to an application, when the flow is redelivered, it tries to

duplicate the data and fails because the data already exists. If the collaboration processing the flow is transactional, a rollback occurs. If the flow is not transactional, it is moved to the resubmission queue. See on page 0 “Resolving failed flows” on page 138 for more information on submitting a flow that fails to process.

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

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

## Changing the InterChange Server Express and database passwords

Password encryption provides a measure of security for protecting the Business Integration Express for Item Sync 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 when the `repos_copy` command is used. See “Using Repos\_Copy” on page 64.

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

“Changing the InterChange Server Express password”

“Changing the database passwords” on page 39

### Changing the InterChange Server Express password

To change the password for InterChange Server:

1. Open System Manager.
2. Right-click the server under Server Instances in the InterChange Servers section of System Manager, then select Change Password. The Change InterChange Server Express Password dialog box appears.
3. Enter the current password in the Old Password field.
4. Enter a new password in the New Password field.
5. Reenter the new password in the Confirm Password field.
6. Click OK.

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



**Attention:** The InterChange Server 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 will not start.

## Changing the database passwords

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

To change the database passwords:

1. Under Server Instances, right-click the server under whose database password you want to change, then select Edit Configuration. The upper-right section of the window changes to an editing tool in which many system properties can be changed.
2. Click the Database tab to access the database configuration properties. The Server Property and Configuration window for database properties appears (see Figure 17).

The screenshot shows the 'InterChangeServer' window with the 'Database' tab selected. The window is divided into four main sections: Database Connectivity, Event Management, Transactions, and Repository. Each section contains a set of configuration fields. In the 'Event Management', 'Transactions', and 'Repository' sections, the 'Password' field is highlighted with a 'Change' button next to it. The 'Database Connectivity' section includes fields for Database Driver (MS SQL Server), Max Connections (50), Max pools (10), Idle Timeout (2 Min), Max deadlock retry (5), and Deadlock retry interval (20 Second(s)). The 'Event Management', 'Transactions', and 'Repository' sections each have fields for Host name (cwe1u48), Database (cwrepos), Login (ics), Password (masked with asterisks), Max Connections (Unlimited), and Port number (1433).

Figure 17. Database tab of Edit Configuration window

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

4. Click OK.

---

## System Manager tasks

This section describes system administration tasks of System Manager, such as starting up, shutting down, and refreshing. For detailed information about using System Manager while implementing the IBM WebSphere ICS system, refer to the *System Implementation Guide*. This section covers the following topics:

“Starting System Manager”

“Shutting down System Manager”

“Refreshing System Manager and updating components” on page 41

“Configuring system-wide flow control” on page 41

## Starting System Manager

To start System Manager, click Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration Toolset > Administrative > System Manager. [WHAT ARE THE NEW SHORTCUT NAMES? THROUGHOUT.] The System Manager perspective of the IBM WebSphere Studio Workbench appears (see Figure 18).




Figure 18. System Manager

## Shutting down System Manager

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

**Note:** Be sure to shut down any InterChange Servers before shutting down System Manager. For instructions on shutting down InterChange Servers, see “Shutting down InterChange Server Express” on page 37.

## Refreshing System Manager and updating components

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

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

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

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

## Configuring system-wide flow control

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control for individual components, see “Configuring flow control for connectors” on page 51 or “Configuring flow control for collaboration objects” on page 56.

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

To monitor how flow control is working in the system, you can use the default Flow Control monitor provided as part of the Web-based System Monitor, or you can create new Flow Control monitors for individual connectors or collaboration objects. For more information about monitoring flow control, see the description of the default Flow Control monitor in “Using default monitors” on page 11, and the instructions on creating new monitors in “Creating new monitors using the Monitor Definition Wizard” on page 18. After you have all the flow control monitors created, refer to the following sections to begin monitoring flow control: “Using the Web-based System Monitor” on page 11 and “Logging on to the Web-based System Monitor” on page 22. “Operating collaboration objects” on page 53

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

1. In System Manager, right-click the server under Server Instances for which you want to configure flow control, then select Edit Configuration. The upper-right section of the window changes to an editing tool in which many system properties can be changed.
2. Click the Misc tab. A dialog box appears with a Flow Control section (see Figure 19).

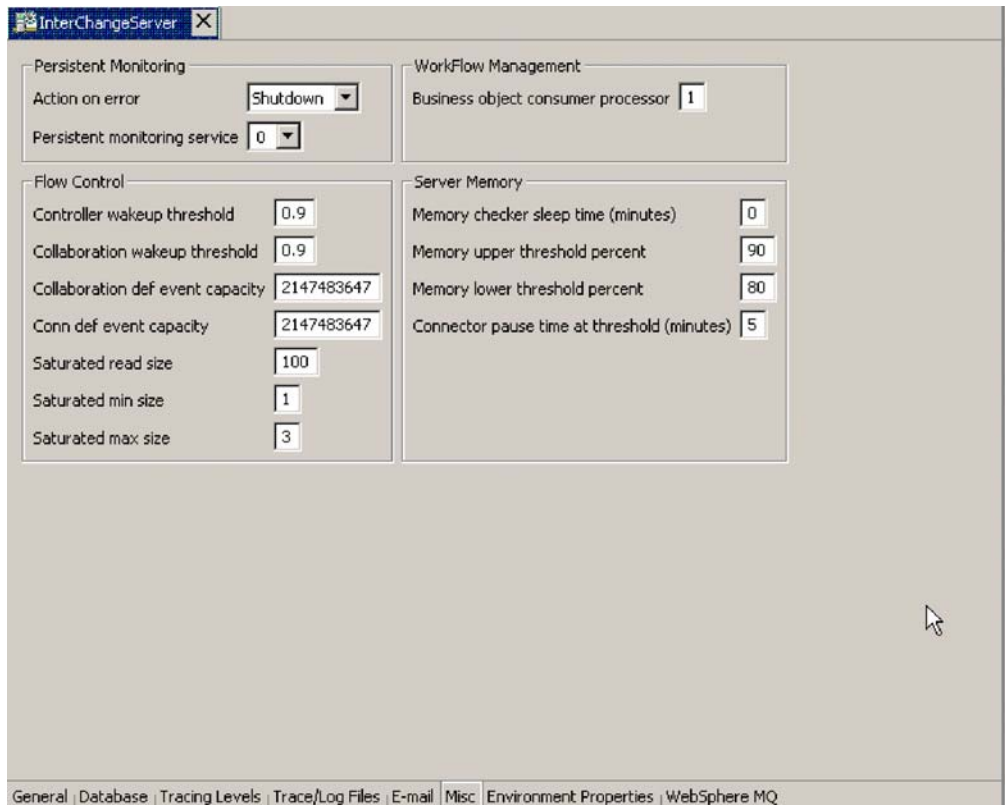


Figure 19. Edit configuration tool, Misc tab

3. In the Flow Control section, enter 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 ICS configuration by selecting Save <ServerName> from the File menu of System Manager.
5. Restart InterChange Server.

---

## Operating connectors

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

You can start, pause, stop, and shut down connectors from the Web-based System Monitor.

This section covers the following topics:

“Connector states” on page 43

“Starting, stopping, and pausing connectors” on page 44

“Configuring flow control for connectors” on page 51

### Connector states

You can view the state of a connector by logging on to the Web-based System Monitor and opening a view that contains connector states.

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

#### Web-based System Monitor

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

1. If the System Overview view is not displayed, click the System Overview link under Views in the left pane of the Web page. The System Overview Monitor appears (see Figure 11 on page 24) in the body of the Web page.

When the product is installed, the default view is set to System Overview, and the default monitor contained in that view is set to System Overview. These defaults can be changed to suit your monitoring needs. See “Using Views” on page 25 for instructions.

2. Click the triangle next to the name of the server to reveal a list of components on the system.
3. Click the triangle next to a running collaboration to reveal its associated connectors (see Figure 20).

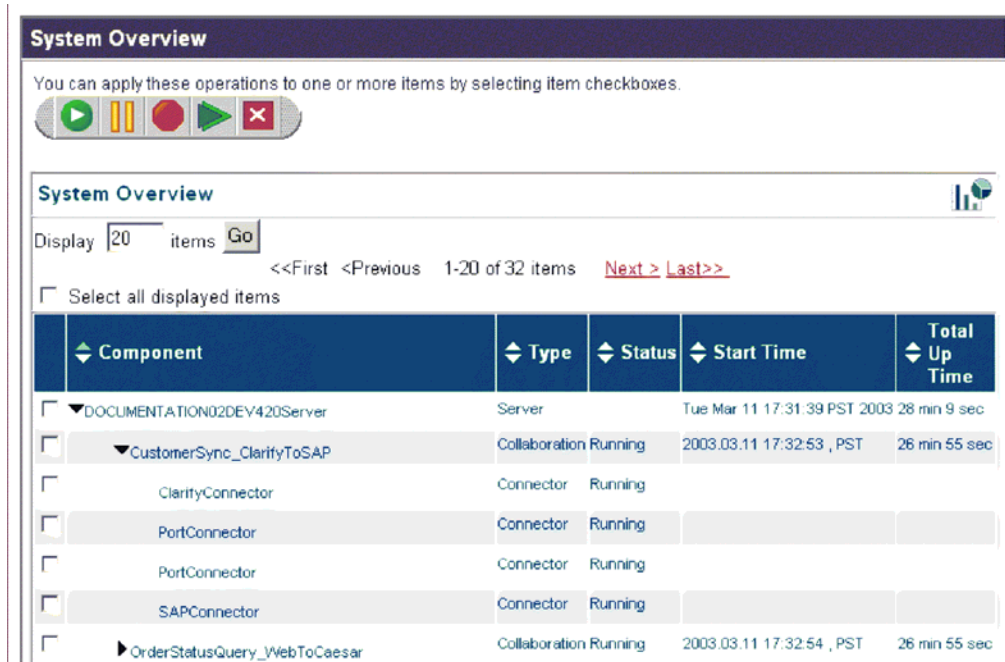


Figure 20. Web-based System Monitor, System Overview displaying connector status

## Starting, stopping, and pausing connectors

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

“Initializing a connector” on page 45

“Running, stopping, and pausing connectors using System Monitor” on page 45

“Commands for changing connector states” on page 45

“Manually starting a connector” on page 46

“Shutting down a connector” on page 47

“Restarting a connector” on page 48

“Setting Automatic and remote restart for a connector” on page 48

“Using OAD as a Windows service” on page 51

“Specifying a log file for use with OAD” on page 51

## Initializing a connector

The first time you start a connector, it must be initialized. Initializing a connector requires that you start it manually. For instructions on manually starting a connector, see “Starting a connector manually on Windows” on page 46.

If the connector does not start, check to make sure that the command line to start it includes the current InterChange Server Express name.

## Running, stopping, and pausing connectors using System Monitor

[OK?] After the connector has been initialized, you can start, stop, and pause it using Web-based System Monitor from System Manager:

### Web-based System Monitor

1. While viewing the System Overview view (see Figure 20 on page 44), select a connector by placing check in the box to its left.
2. Select the Start, Pause, or Stop icon from the icon group in the upper-left corner of the view (see Figure 21).



Figure 21. Web-based System Monitor, icons for starting, pausing, restarting, or shutting down components

## Commands for changing connector states

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

Start <i>NameConnector</i>	Starts the connector for the application <i>Name</i> if it is paused or stopped. Connectors poll the application and connector controllers read the persistent queue. Flows are processed.
Pause <i>NameConnector</i>	Pauses the connector for the application <i>Name</i> if it is running or stopped. Connectors stop polling the application and connector controllers stop reading the persistent queue. Flows are not processed.
Stop <i>NameConnector</i>	Stops the connector for the application <i>Name</i> if it is running or paused. Connectors stop polling the application and fail requests with an exception message. Connector controllers stop reading the persistent queue. Flows and requests are not processed.
Shut Down <i>NameConnector</i>	Shuts down the connector for the application <i>Name</i> . The connector’s process is stopped.
Boot Up Connector Agent	Restarts the connector for the application <i>Name</i> . This action is available only if you have set the <code>OADAutoRestartAgent</code> property of the connector to <code>True</code> . See “Setting Automatic and remote restart for a connector” on page 48.



## Manually starting a connector

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

You can start the connector in several ways:

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

### Java connector

```
start_Sap ConnectorName InterChangeServerName
```

### C++ connector

```
start_connector ConnectorName InterChangeServerName
```

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

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

You can customize the startup for each connector by modifying the connector shortcut or the `start_connector.bat` file. Use the connector startup parameters listed in Table 6 to customize the startup of a connector.

Table 6. Connector Startup Parameters

Parameter	Function
-c <i>configFile</i>	Name of the configuration file to be used during startup. If the filename specifies a relative path, the startup script looks for the file in the directory where the product is installed. This parameter is required only to use a local connector configuration file. If you are not using a local configuration file, enter the name of the configuration file used by 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.



Table 6. Connector Startup Parameters (continued)

Parameter	Function
-d	Specifies the name of the C++ connector's library file, which is a dynamic link library (DLL). This DLL name does not include the .dll file extension. The startup script specifies this option for all C++ connectors.
-f <i>pollFrequency</i>	<p>Poll frequency is the number of milliseconds between polling actions.</p> <ul style="list-style-type: none"> <li>To specify the number of milliseconds, provide a value for <i>pollFrequency</i>.</li> <li>To cause the connector to poll only when you type the value <i>p</i> in the connector's Command Prompt window, specify the <i>-fkey</i> option.</li> <li>If a connector is configured to processes only business object requests and not application events, polling is unnecessary; you can disable polling by specifying <i>-fno</i>.</li> </ul> <p>The value of this parameter overrides any repository definitions. You can specify either <i>-fkey</i> or <i>-fno</i>, but not both.</p>
-j	Specifies that the connector is written in Java. This parameter is optional if you specify <i>-l className</i> .
-l <i>className</i>	Specifies the name of the Java connector's global class, which is an extension of the connector base class. The startup script specifies this option for all Java connectors.
-n <i>connectorName</i>	Specifies the name of the connector to start.
-p <i>password</i>	Specifies the password that the connector uses to access InterChange Server.
-s <i>serverName</i>	Specifies the name of the InterChange Server. This parameter is required. The name is case-sensitive.
-t	Turns on the connector property <code>SingleThreadAppCalls</code> . This property guarantees that all calls the connector framework makes to the application-specific connector code are with one event-triggered flow. The default value is <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.sh` script, which handles the actual connector management for the connector. The IBM WebSphere ICS product provides a `start_connector.sh` script for each connector it delivers. Shutting down a connector stops the connector's processes. Before shutting down a connector, pause or stop each collaboration object that uses the connector (the collaboration must be configured to pause; see the collaboration documentation for details on how to do this). If the "Pause when critical error occurs" property has been set for a collaboration in the Collaboration General Properties window, the collaboration pauses automatically when a critical error occurs. The latest unprocessed events of such collaborations are then moved to the event submission queue.

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

- If you have not enabled automatic restart, when you perform a shutdown action the effect is “permanent”—that is, the connector shuts down and will not restart until and unless you restart it manually at the command line or with a batch file.
- If you have enabled automatic restart, the shutdown action is temporary, and you can restart the connector by using the Boot Up Connector Agent action in System View.

For instructions on enabling or disabling automatic restart, see “Setting Automatic and remote restart for a connector” on page 48.

To shut down a connector:

#### UNIX

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

#### Windows

Use System Monitor to shut down a connector.

To shut down a connector using Web-based System Monitor:

1. From the System Overview view, select the collaboration object of the connector you want to shut down by placing a check in the box to its left, then click the Pause icon from the upper-left corner of the view (see Figure 21 on page 45). Do this for each collaboration associated with the connector.
2. Select the connector you want to shut down by placing check in the box to its left, then click the Shutdown icon from the upper-left corner of the view (see Figure 21 on page 45).

### Restarting a connector

This action is used to restart the connector after you have used the Shut Down Connector action the Web-based System Monitor. This action is available only if you have enabled automatic and remote restart for the connector (see “Setting Automatic and remote restart for a connector” on page 48).

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

#### Web-based System Monitor

1. From the System Overview view (see Figure 20 on page 44), place a check in the box to the left of the connector you want to restart.
2. Click the Restart Agent icon from the upper-left corner of the view (see Figure 21 on page 45).

### Setting Automatic and remote restart for a connector

You can enable a connector to attempt to restart automatically after it has shut down abnormally and to be capable of a remote restart from System Monitor.

This feature has the following restrictions and prerequisites:

- It cannot be used for a Secure Connect installation of a connector
- For JMS, this feature can be enabled only in an intranet environment

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

Before enabling remote and automatic restart for a connector, check the setting of the OSAGENT\_PORT value, and edit the JVM parameters in the configuration file for the connector.

**Checking the Osagent port value:** By default, the environment variable OSAGENT\_PORT value is set to 14000 when you install the Borland VisiBroker software. If you keep the default value, you do not need to change the osagent port value setting. However, if you change the OSAGENT\_PORT environment variable from the default--for example, if you change the value from 14000 to 14010--you must make the following changes:

#### UNIX

Set the environment variable OSAGENT\_PORT to the port number you intend to use for the osagent and export the value; for example:

```
set OSAGENT_PORT=14010
export OSAGENT_PORT
```

#### Windows

Set the OSAGENT\_PORT environment variable, and set the vbroker.agent.port value in the agent configuration file to the same value as the OSAGENT\_PORT environment variable.

**Editing JVM parameters:** Before enabling remote and automatic restart for a connector, you must edit the JVM<connector\_name> parameter in the configuration file for the connector.

Use an XML editor to open the InterchangeSystem.cfg file, or a text editor to open the local connector configuration file if you are using one, and edit the file to set values for parameters under the JVM <connector\_name> section. Set the values to exactly match the corresponding values in the batch or script file for that connector.

The JVM parameters and their corresponding Java VM options are:

```
MIN_HEAP_SIZE (-Xms)
MAX_HEAP_SIZE (-Xmx)
MAX_NATIVE_STACK_SIZE(-Xss)
```

For example, the JVM section of the configuration file might set the following values:

```
[JVM SAPConnector]
MIN_HEAP_SIZE=256m
MAX_HEAP_SIZE=512m
MAX_NATIVE_STACK_SIZE=1m
```

**Enabling OAD for connectors:** To enable remote starting and automatic restart to be used by a connector for the first time, perform these steps:

1. Start the OAD daemon, using one of these commands:

**UNIX**

```
$PRODUCTDIR/bin/start_ActivationDaemon.sh
```

**Windows**

```
PRODUCTDIR\bin\start_ActivationDaemon.bat
```

2. Start the connector.
3. Start ICS.
4. From System Manager, double-click the connector. This opens Connector Configurator.
5. Set the following standard properties in the Standard Properties tab. To enable OAD for the connector, set the OADAutoRestartAgent value to True:

*Table 7. Configuring standard properties in Connector Configurator*

Name	Possible Values	Description	Default Values
OADAutoRestartAgent	true or false	If set to true, the OAD will automatically attempt to restart the connector after an abnormal shutdown. It can also be used to start the connector agent remotely. This value is dynamic.	false
OADAutoRestartAgent	true or false	If set to true, the OAD will automatically attempt to restart the connector after an abnormal shutdown. It can also be used to start the connector agent remotely. This value is dynamic.	false
OADAutoRestartAgent	true or false	If set to true, the OAD will automatically attempt to restart the connector after an abnormal shutdown. It can also be used to start the connector agent remotely. This value is dynamic.	false

The steps described above register one connector with the OAD. When you have performed these steps, OAD becomes active for that connector, and will automatically restart the connector until and unless you deactivate the OAD feature by changing the OADAutoRestartAgent value to False.

**Toggle Automatic and Remote Restart:** To toggle the automatic and remote restart feature on and off, change the value of the OADAutoRestartAgent property in the Connector Configurator window for the connector. The change is dynamic: When you set the value to False, automatic restart is disabled. When you set the value to True, automatic restart is enabled. You do not need to restart ICS for the change to take effect.

**Note:** The connector must be running in order for the automatic and remote restart to be toggled on and off. If you toggle the feature off, and the connector subsequently goes down, you will not be able to use automatic or remote start until you again perform the steps described under “Enabling OAD for connectors” on page 49.

**If OAD Terminates Abnormally:** If the connector shuts down abnormally and the OAD also terminates abnormally before it is able to restart the connector, the connector will not be automatically restarted when you restart OAD. If this occurs, you will need to restart the connector.

### Using OAD as a Windows service

If you are configuring a connector for automatic restart or remote starting using OAD, do not install the connector to run as a Windows service. Instead, install OAD to run as a Windows service. When the Windows system starts, OAD will automatically start as a Windows service; when ICS restarts, it will start the connector through OAD.

To install OAD as a Windows service:

1. In the Windows services tool, locate the “VisiBroker Activation Daemon” entry and change its startup value to Disabled.
2. In a DOS window, obtain a command prompt for the directory that contains your `start_ActivationDaemon.bat` file; the file should reside in your `ProductDir\bin` directory.
3. At the command prompt, enter this command:  
`start_ActivationDaemon.bat -install cw_OAD`  
where `cw_OAD` is the name you choose for your OAD Windows service name.
4. [IS THERE STILL SUCH A TOOL? IF NOT, HOW DO YOU DO THIS?]Run the Windows Service Setup tool. Select “InterChange Server” in the Service Component field, and in the Service Dependencies field enter the name for your OAD Windows service.

After installing OAD as a Windows service, you can use the Services tool in Windows to start and stop the daemon.

**Note:** When you run OAD as a Windows service for restarting the connector, no STDOUT messages for the connector are displayed--instead there will be an empty popup window for each connector when the connector starts. To remove the empty window, go to Control Panel in Windows, choose Services, and select the entry for `cw_OAD` display. Choose Startup. In the Services popup window, deselect the “Allow Service to Interact with Desktop” checkbox, and choose OK.

### Specifying a log file for use with OAD

When OAD activates a process, it redirects STDOUT/STDERR messages to its own window. If you are running multiple connectors using OAD, this can result in a confusing display of messages. To avoid this, in the `LOG_FILE` parameter of the connector configuration file, instead of using STDOUT, specify a path and file name for the output. This will direct STDOUT messages for that connector to the specified file, rather than to the OAD window display.

## Configuring flow control for connectors

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If

you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control system-wide, see “Configuring system-wide flow control” on page 41. This section describes how to configure flow control for connectors.

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

To monitor how flow control is working in the system, you can use the default Flow Control monitor provided as part of the Web-based System Monitor, or you can create new Flow Control monitors for individual connectors or collaboration objects. For more information about monitoring flow control, see the description of the default Flow Control monitor in “Using default monitors” on page 11, and the instructions on creating new monitors in “Creating new monitors using the Monitor Definition Wizard” on page 18. After you have all the flow control monitors created, refer to the following sections to begin monitoring flow control: “Using the Web-based System Monitor” on page 11 and “Logging on to the Web-based System Monitor” on page 22.

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

1. In System Manger, navigate to the connector for which you want to configure flow control, then double-click that connector. Connector Configurator opens (see Figure 22).

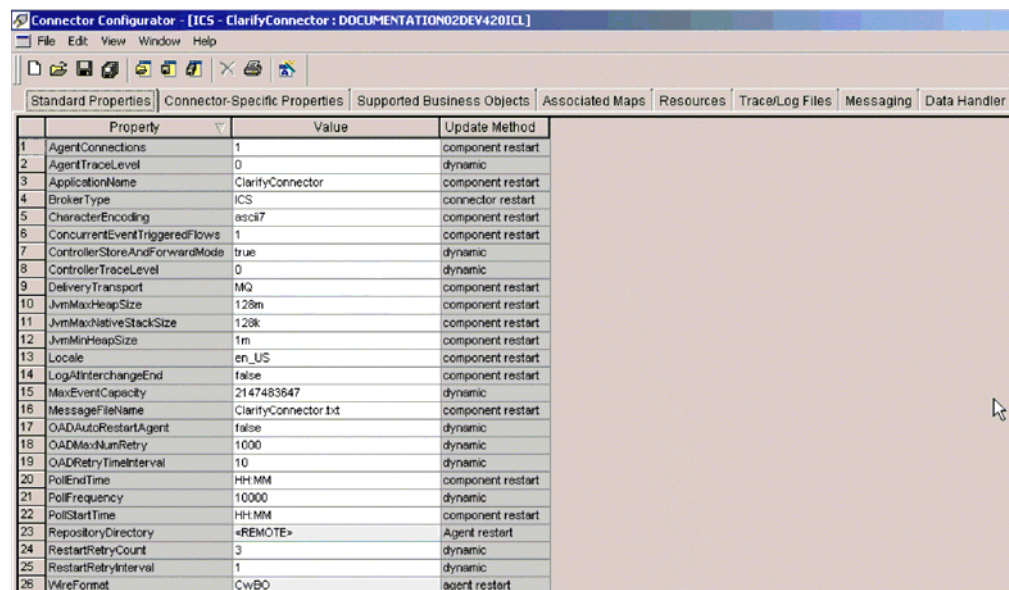


Figure 22. Connector Configurator, Standard Properties tab

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



---

## Operating collaboration objects

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

You can run, pause, stop, and shut down collaboration objects from the Web-based System Monitor.

By default, all connectors and the active collaborations appear.

This section covers the following topics:

“Collaboration object states” on page 53

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

“Configuring collaboration run-time properties” on page 54

### Collaboration object states

You can view the state of a collaboration object by logging on to the Web-based System Monitor and opening a view that contains collaboration object information. To log on to the Web-based System Monitor, follow the instructions in “Logging on to the Web-based System Monitor” on page 22.

The following collaboration object states are viewable from the Web-based System Monitor and the Collaboration Object menu of the System View window:

Start	Starting a collaboration object causes it to subscribe to its triggering business objects and to process them as they arrive. If you stop and then restart InterChange Server, collaboration objects in the Start state automatically start running when InterChange Server Express comes back up.
Pause	<p>Pausing a collaboration prevents it from receiving new flow initiators. The collaboration completes all of the current processing, then enters an idle state.</p> <p>A connector maintains its subscription information; therefore, it continues to send flow initiators to the connector queues. The collaboration processes these when it is resumed.</p> <p>To resume collaboration execution, click Start in the Web-based System Monitor.</p>
Stop	<p>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.</p> <p>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.</p>
Shut Down	Shutting down a collaboration immediately ends processing of current flows. When the collaboration is restarted, the system recovers by processing those flows that were interrupted by the

shut down and recovering those flows waiting in the queue. This recovery is not immediate, so prepare to wait while the system completes the recovery interval.

**Statistics** When you choose Statistics, the Open Collaboration dialog box displays. Enter a collaboration name and choose Open to display statistics for that collaboration object.

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

## Starting, stopping, and pausing collaboration objects

To make a collaboration object functional for the first time, you must first configure it then start it. See “Configuring collaboration run-time properties” on page 54 for more information on configuring collaborations. Depending on which version of System Monitor you are using, you run, stop, and pause collaboration objects in the following way:

### Web-based System Monitor

1. While viewing the System Overview view (see Figure 20 on page 44), select a collaboration object by placing check in the box to its left.
2. Select the Start, Pause, or Stop icon from the icon group in the upper-left corner of the view (see Figure 21 on page 45).

## Configuring collaboration run-time properties

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

“Setting collaboration object general properties” on page 54

“Configuring collaborations to process concurrent event-triggered flows” on page 56

“Configuring flow control for collaboration objects” on page 56

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

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

### Setting collaboration object general properties

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

1. From System Manager, right-click a collaboration object and select Properties.
2. In the Properties dialog box, choose the Collaboration General Properties tab. The following dialog displays:



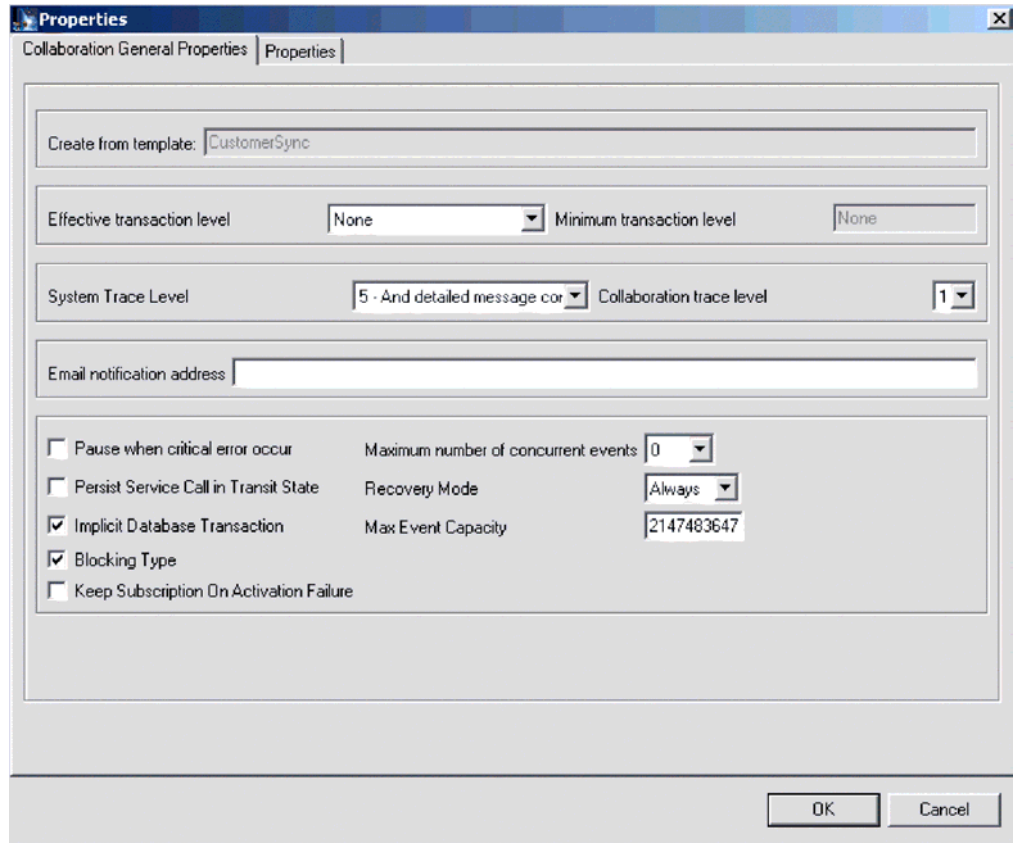


Figure 23. Properties dialog box, Collaboration General Properties tab

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

The dialog enables you to make settings for the following:

- Effective Transaction Level  
See “Setting the Effective Transaction Level and Other General Properties” in the *System Implementation Guide*.
- System Trace Level  
See on page 0. “Configuring logging and tracing” on page 117
- Collaboration Trace Level  
See on page 0. “Configuring the collaboration trace level” on page 123
- Email Notification Address  
See . “Configuring e-mail notification at the collaboration object level” on page 113
- Pause When Critical Error Occurs  
See on page 0 “Critical errors” on page 136.
- Persist Service Call In Transit State  
See ,Chapter 4, “Troubleshooting the system,” on page 105, and see the *Collaboration Development Guide*.
- Implicit Database Transaction  
See “Configuring Transaction Bracketing” in the *System Implementation Guide*.
- Maximum Number of Concurrent Events

See “Configuring collaborations to process concurrent event-triggered flows” on page 56.

- Recovery Mode

See on page 0. “Deferred recovery of collaboration events” on page 134

## Configuring collaborations to process concurrent event-triggered flows

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

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

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

1. From System Manager, right-click the collaboration object that you want to change, then select Properties. The Properties dialog box appears (see Figure 23 on page 55).
2. In the Collaboration General Properties tab, enter a value in the “Maximum number of concurrent events” field.
3. Click OK to save your changes and close the window.
4. Restart the collaboration for changes to take effect.

## Configuring flow control for collaboration objects

Flow control is a configurable service that allows you to manage the flow of connector and collaboration object queues. The parameters for configuring flow control can be configured system-wide or on individual components, or both. If you configure both, the individual component configuration supersedes the system-wide configuration. For instructions on configuring flow control system-wide, see “Configuring system-wide flow control” on page 41. This section describes how to configure flow control for collaboration objects.

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

To monitor how flow control is working in the system, you can use the default Flow Control monitor provided as part of the Web-based System Monitor, or you can create new Flow Control monitors for individual connectors or collaboration objects. For more information about monitoring flow control, see the description of the default Flow Control monitor in “Using default monitors” on page 11, and the instructions on creating new monitors in “Creating new monitors using the Monitor Definition Wizard” on page 18. After you have all the flow control monitors created, refer to the following sections to begin monitoring flow control: “Using the Web-based System Monitor” on page 11 and “Logging on to the Web-based System Monitor” on page 22.

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

1. In System Manger, right-click the collaboration object for which you want to create flow control, then select Properties from the drop-down menu. The Properties dialog box appears (see Figure 23 on page 55).

2. In the Collaboration General Properties tab, edit the value in the Max Event Capacity field to represent the maximum number of events you want queued for a collaboration object. The valid range of values for this property is from 1 to 2147483647.
3. Click OK. The property is changed immediately.

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

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

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

1. From System Manager, right-click the collaboration object whose service call timeout value you want to edit, then click Properties. The Properties dialog box appears.
2. From the Properties tab, locate the property that represents the service call timeout value you want to change, then click in the value field. When the property becomes highlighted, the value can be edited.

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

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

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

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

---

## **Operating maps**

You can start and stop maps from the Web-based System Monitor.

This section covers the following topics:

"Map states"

"Starting and stopping maps" on page 58

### **Map states**

You can view the state of a map by logging on to the Web-based System Monitor.

#### **Web-based System Monitor**

To see the state of connectors using the Web-based System Monitor, do the following: If the System Overview view is displayed, click the Maps and Relationships link under Views in the left pane of the Web page. The Map Status

and Relationship Status monitors appear in the body of the Web page. When the product is installed, the default view is set to System Overview, and the default monitor contained in that view is set to System Overview. These defaults can be changed to suit your monitoring needs.

## Starting and stopping maps

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

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

### Map activation

For a map to be executable, it must first be activated. Map Designer automatically starts a map when it successfully compiles the map. However, other changes to the map might require that you explicitly stop and restart the map for the change to take effect. See Appendix, “Requirements for restarting IBM WebSphere Business Integration system components,” on page 155 to find out what changes require maps to be stopped and restarted.

### Starting and stopping maps

Depending on which version of System Monitor you are using, you start and stop maps in the following way:

#### Web-based System Monitor

1. If the System Overview view is displayed, click the Maps and Relationships link under Views in the left pane of the Web page. The Map Status and Relationship Status monitors appear (see Figure 11 on page 24) in the body of the Web page.
2. Select the Start or Stop icon from the icon group in the upper-left corner of the view (see Figure 21 on page 45).

---

## Operating relationships

You can start and stop relationships from the Web-based System Monitor.

This section covers the following topics:

“Relationship states” on page 58

“Starting and stopping relationships” on page 59

### Relationship states

You can view the state of a relationship by logging on to the Web-based System Monitor and opening a view that contains relationship status. To log on to the Web-based System Monitor, follow the instructions in “Logging on to the Web-based System Monitor” on page 22.

Table 8 lists the relationship states represented by the System View traffic light display color and describes what actions can be performed during that state.

*Table 8. Relationship States*

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

## Starting and stopping relationships

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

When you expand the Relationships folder in System Manager, two subfolders appear: Dynamic and Static.

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

This section describes the following topics:

“Relationship activation” on page 59

“Starting and stopping relationships” on page 60

“Relationship table caching” on page 60

### Relationship activation

For a relationship to be executable, it must be activated. However, you cannot modify a relationship when it is active. Therefore, you must stop the relationship, make the change to the relationship, and then restart the relationship. See Appendix , “Requirements for restarting IBM WebSphere Business Integration system components,” on page 155 to find out what changes require relationships to be restarted.

## Starting and stopping relationships

Depending on which version of System Monitor you are using, you start and stop relationships in the following way:

### Web-based System Monitor

1. From the System View window (see Figure 11 on page 24), select Show Maps and Relationships from the View drop-down menu. The Maps and Relationships columns appear next to Collaborations and Connectors.
2. Right-click a relationship, then select either the Start<name\_of\_relationship> or Stop <name\_of\_relationship> option.

### Relationship table caching

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

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

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

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

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

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

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

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

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

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

3. Right-click the static relationship to determine its current cached state. If the Cached option appears with a check mark to the left, caching for that relationship is currently enabled. Therefore, the Reload option is enabled.
4. Choose Reload from the context menu to reload the static relationship's tables.

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

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

```
RELATIONSHIP.CACHING=5
```

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

---

## Backing up IBM WebSphere ICS components

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

This section covers the following topics:

“Planning a Backup Schedule” on page 61

“Backing Up Components” on page 62

### Planning a Backup Schedule

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

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

- Static data rarely changes and should be backed up only when changed. For example, static configuration data stored in the IBM WebSphere ICS repository needs backing up only when it has changed. Static data should also be backed up before any planned reinstallations or upgrades to the system. Following is a partial list of static data in the IBM WebSphere ICS system:
  - IBM WebSphere ICS repository (except for relationship tables)
  - Custom collaborations components, such as Java class files (.class), and message files (.msg)
  - Custom connectors
  - Map components, including: map design files and Java class files (.class).



- Dynamic data constantly changes and should be backed up on a regular basis. For example, the relationship tables maintain the instance data for relationship definitions. Since relationship instance data is maintained continuously, regularly back up this data as well as application data.

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

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

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

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

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

## Backing Up Components

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

“Backing Up Relationship Tables” on page 62

“Backing Up and Loading the Repository” on page 63

“Backing up System Installation Files” on page 63

“Backing Up Collaboration Class Files” on page 63

“Backing Up Archive Tables” on page 63

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

### Backing Up Relationship Tables

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



relationship tables. For more information on bringing the system to a quiescent state, see “Shutting down InterChange Server Express” on page 37.

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

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

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

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

### **Backing Up and Loading the Repository**

Repository tables are backed up using the `repos_copy` command. For more information on this command, see “Using Repos\_Copy” on page 64. Back up the repository whenever it is modified and before and after performing a reinstallation or an IBM WebSphere ICS software upgrade. The IBM WebSphere ICS system does not need to be in a quiescent state when backing up the repository.

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

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

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

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

### **Backing up System Installation Files**

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

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

### **Backing Up Collaboration Class Files**

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

### **Backing Up Archive Tables**

Some applications have archive tables. Back up archive tables using the standard database utility for the database in which they reside. The archive tables are part of

the IBM WebSphere ICS system, but typically reside in the application's database. Back up the archive tables on a regular basis. Data in the archive table represents all of the events that have passed from the application to the IBM WebSphere ICS system. These events can be used to "resynchronize" the application and the IBM WebSphere ICS cross-reference tables. For more information, see on page 0. "Maintaining event archives" on page 132

---

## Using Repos\_Copy

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

To run `repos_copy`, enter the command 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` will not work.

**Important:** When `repos_copy` deploys components to the repository, it deploys them to the repository *only*. It does *not* deploy them to any in-memory tables of business object definitions. For instance, connectors load business object definitions from the repository into their memory space when they start. If you deploy a business object definition to the repository to update it, you must restart the connector agent so that it loads the modified business object definition into memory. You must therefore stop and restart InterChange Server 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 64
- "Repos\_copy usage scenarios" on page 71
- "Locale for repos\_copy files" on page 76

### Repos\_copy syntax

Table 9 on page 65 describes the options of `repos_copy` and their arguments, and shows the correct case usage for the options and the lack of spacing between the option and its argument. The syntax shows that the options between curly braces ({} ) represent a set of options that are required. If you do not specify the `-u`, `-p`, `-i`, `-o`, or `-s` options at the command line, then `repos_copy` prompts you for them. If you do not specify them when prompted, `repos_copy` does not execute. Options enclosed in brackets ([]) are optional.

**Note:** Some new arguments have been added in release 4.2, and some arguments from the previous release have been removed. For a list of these arguments, see "New arguments in release 4.2" on page 70 and "Arguments removed in release 4.2" on page 71.

```
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 9. *Repos\_copy* command options

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

Table 9. *Repos\_copy* command options (continued)

Option	Description
<code>-doEntityType:Entity[+EntityType:Entity2]</code>	<p>Specifies the entities to be deleted from the repository. See Table 10 on page 70 for the list of entity types and keywords. If the object has no referents—other components that depend on it—then the deletion takes place. If the object has referents, then the deletion fails and a message is displayed. The behavior is the same in both design mode and production mode. For more information about starting the server in design mode or production mode, see the <i>System Implementation Guide</i>.</p>
<code>-eEntityType:Entity1[+EntityType:Entity2...]</code>	<p>Exports one or more referenced first-class entities. A first-class entity is a business object, collaboration object, collaboration template, connector, database connection pool, map, or relationship. You identify the entity to load or unload by specifying one of the keywords in Table 10 on page 70.</p> <p>Follow the <i>EntityType</i> keyword with a colon (:) and the name of the entity. Use the “+” to specify more than one entity. When combined with the <code>-o</code> option, the <code>-e</code> option unloads the data to an output file.</p>
<code>-fentityFile</code>	<p>This option is similar to the <code>-e</code> option except that the names of the entities to be imported are stored in a file. The file should contain references to the entities, with the following conditions:</p> <ul style="list-style-type: none"> <li>• The entity names must follow after the proper entity type keyword. The entity types and their keywords are listed in Table 10 on page 70.</li> <li>• A colon must separate the entity type from the entity name.</li> <li>• There must be a new line separating each entity reference.</li> </ul> <p>When combined with the <code>-o</code> option, this option exports the components to a package.</p>

Table 9. *Repos\_copy* command options (continued)

Option	Description
<i>-ifilename</i>	<p>Deploys the specified package file to the repository. If you omit the input file name value, the command interactively prompts you to enter the name of the input file. The file can be either a .jar file containing objects in XML format, or a file in text format from a release prior to 4.2.0.</p> <p>The .jar 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>Overrides the default behavior of <i>repos_copy</i> when it finds a Mercator map in the package file it is loading. By default, <i>repos_copy</i> exits if it encounters a Mercator map. If you use the <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>Returns the mode of the server.</p>
<i>-ncencoding</i>	<p>Specifies the character encoding when importing a text-based repository file from releases prior to 4.2.0.</p> <p>For a list of valid character encodings, see the Java documentation about the String class.</p>

Table 9. *Repos\_copy* command options (continued)

Option	Description
<i>-outfilename</i>	Exports the components in the repository to the specified package file. You must specify the name of the package file. If the file already exists then <i>repos_copy</i> prompts you to overwrite it or not. The output file is in .jar format, and contains the component definitions in XML format, as well as .java source files for components that have them. This option cannot be combined with the <i>-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 .jar 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>	This option is similar to the <i>-r</i> option; it allows you to import relationship definitions and not create the runtime schemas for any of them.
<i>-rrelationshipName1[:relationshipName2]</i>	Loads the named relationship definition(s) into the repository without creating its runtime schema.
<i>-sserverName</i>	Specifies 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>	This option prints a list of components in the server repository (they are identified as "artifacts" rather than as components in the output). The output is in XML format. You can combine this option with the <i>-o</i> option to print the output to a file rather than the console.

Table 9. Repos\_copy command options (continued)

Option	Description
<code>-username</code>	Specifies the user name to log in to InterChange Server. If no user name is specified, repos_copy prompts for a user name.
<code>-v</code>	Prints the version number of the program that the repos_copy utility executes.
<code>-vp</code>	This option validates a package file. The server validates packages against the repository and makes sure that the dependencies among the components in the package are resolved. If the validation is not successful, repos_copy prints a list of the missing dependencies. This option does not make any changes to the repository; it just validates the package file. When using the <code>-vp</code> option you must also use the <code>-i</code> option to specify the package file to be validated.
<code>-vr</code>	This option validates the repository. The output message indicates whether the validation is successful or not. If the validation is not successful, repos_copy prints a list of the missing dependencies.
<code>-wi</code>	When this option is specified, repos_copy does not display any warnings that occur during the compilation of collaboration templates or maps. Only errors that occur during compilation are displayed. This allows the user to ignore warnings about deprecated methods, for instance.
<code>-xCompileAll</code>	Compiles all collaboration templates and maps in the repository. Valid only for collaboration templates and maps created using release 4.2 or later.
<code>-xCompileAllCollabs</code>	Compiles all collaboration templates in the repository. Valid only for templates created using release 4.2 or later.
<code>-xCompileAllMaps</code>	Compiles all maps in the repository. Valid only for maps created using release 4.2 or later.

Table 9. Repos\_copy command options (continued)

Option	Description
<code>-xCompileCollab:collabTemplateName[+collabTemplateName]</code>	Compiles the specified collaboration templates in the repository. Valid only for templates created using release 4.2 or later.
<code>-xCompileMap:nativeMapName[+nativeMapName]</code>	Compiles the specified maps in the repository. Valid only for maps created using release 4.2 or later.
<code>-xCompilePackage</code>	This option automatically compiles the package being deployed to the server. Since the production-mode server automatically compiles all packages, this option applies only to design-mode servers. For a full description of InterChange Server Express modes, see the System Implementation Guide. <b>Note:</b> This option works only if you are deploying components from release 4.2. If the components are from a prior release, this option will be ignored.

Table 10. Keywords for different entity types

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

## New arguments in release 4.2

The following list contains all of the new options provided with repos\_copy in release 4.2:

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



## Arguments removed in release 4.2

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

Table 11. Arguments removed from repos copy

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

## Repos\_copy usage scenarios

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

- “Printing the `repos_copy` command”
- “Validating a package”
- “Validating a package”
- “Deploying a package to the repository” on page 72
- “Validating the repository” on page 73
- “Deleting components from the repository” on page 74
- “Exporting components to a package” on page 75
- “Printing a list of components in the repository” on page 76

### 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]
```

### Validating a package

You can validate a package of components before deploying the package to a server. This is very useful because if you deploy a package to a production-mode server all the dependencies must be resolved or the deployment will fail. You

cannot validate a user project or integration component library in System Manager to make sure that the dependencies are satisfied, so the only way to find out if a package is valid when deploying with System Manager is to attempt the deployment and use the error information when it fails to resolve the dependencies. If there are many components in the package, this can be a very time-consuming process.

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

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

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

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

### Deploying a package to the repository

The `-i` option allows you to deploy a package of components to the repository. If you do not specify the name of the package file then you are prompted to enter it.

The following example shows a file named `WebSphereICS420DEVServer.jar` being deployed to a repository:

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

**Working with duplicate components during deployment:** Commonly there will be components with the same name in the package file as there are in the repository. In this case you must decide whether or not you want to replace the components in the repository with those in the package file. The `-ai` option specifies that duplicate components should not be loaded into the repository:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-iCustomer.jar -ai
```

If you want to replace all the duplicate components in the repository, use the `-ar` option as in the following example:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-iCustomerSyncInterface.jar -ar
```

You can use the `-arp` option to interactively replace duplicate components in the repository. This lets you decide for each individual duplicate component whether it should be replaced or not.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-iCustomerSyncInterface.jar -arp
```

**Note:** The `-ar` and `-arp` options only work with release 4.2.0 or later.

**Compiling and creating schemas:** For maps and collaboration to execute at runtime, the maps and collaboration templates defined in the repository must be compiled. For relationships to function properly at runtime, their schemas must be created.

When you deploy components to a server running in production mode, all templates are automatically compiled and all relationship schemas are created. For the deployment to succeed, then, the code of the map and collaboration templates

must be valid and InterChange Server 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 instance, if you are deploying a package from one environment to another and did not change the properties of the relationships to use the database resources in the new environment then you will not want the schemas created until after you have changed the relevant properties. The following example uses the `-r*` option to not create schemas for all of the relationships in the package being deployed:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-iWebSphereICS420DEVServer.jar -xCompilePackage -r*
```

**Note:** You can use the `-r` option without the asterisk to specify the names of individual relationships whose schemas should not be created. For instance, `-rCustomer:Order` would not create schemas for the Customer and Order relationships, but would still create schemas for any other relationships in the package being deployed.

**Important:** Although there are options to compile maps and collaboration templates after deployment, there is no way to either through `repos_copy` or System Manager to create the schema for a relationship other than during deployment. So, if you chose not to create the schema for a relationship during deployment because you needed to change the database settings, then you need to re-deploy the relationship afterwards and allow `repos_copy` to create the schema for the relationship.

### Validating the repository

The repository must be in a valid state for a server instance to start in production mode. The reason for this is that ultimately the repository must be valid for the server to process flows successfully. Use the `-vr` option to validate a server repository, as in the example below:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull -vr
```

If the server is valid then `repos_copy` writes the following output to the console:

```
Validation Succeeded.All Dependencies Resolved.
```

If the repository is not valid then `repos_copy` prints a list of the dependencies that must be resolved.

### Compiling components in the repository

If you deployed maps or collaboration templates to the repository and did not compile them during deployment, you can use `repos_copy` to compile them afterwards. This can be useful in situations where there are many components to deploy because deployment can take a long time and compiling can make the operation take even longer. Waiting until after the deployment has succeeded to do

the compilation task can reduce the risk of spending an even greater amount of time migrating the environment if an error occurs.

The following example shows the use of the `-xCompileAll` option to compile all maps and collaboration templates in the 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
```

### Deleting components from the repository

There are several options provided by `repos_copy` for deleting components in the repository. You can delete the entire repository, individual components, and individual components as well as any components that reference them.

**Note:** Components must be inactive for you to delete them. If you delete a single component then you must deactivate it first or the delete operation will fail. If you want to delete a component and all the components that reference it, you must deactivate not only the single component, but all those that reference it as well. You can delete the entire repository while the components are in an active state. Use System Monitor or web-based System Monitor to manage the states of components. System Monitor and web-based System Monitor are described in the *User Guide*.

**Deleting the entire repository:** Use the `-d` option to delete all of the components in the repository. The following example shows the syntax:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin  
-pnull -d
```

`Repos_copy` presents a prompt asking if you want to delete the entire repository or not.

**Deleting components without referents:** If a component does not have any referents—other components that reference it and require it to exist in order to perform their function in the system—then you can delete the individual component.

Use the `-do` option followed by the entity type, a colon, and the name of the component. The entity types are listed in Table 10 on page 70. The following example deletes the relationship named **Customer**:

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

**Deleting components with referents:** If a component does have referents—other components that reference it and require it to exist in order to perform their function in the system—then you can only delete the component if the server is running in design-mode, and by using certain options.

*Forcing a delete in spite of references:* If a component has referents, `repos_copy` will not let you delete it with the `-do` option. You must use the `-dfo` option to force deletion of a component with referents. Forcing deletion of a component that has referents will leave the repository in an inconsistent state, and a server running in production mode does not permit that, so this option only works with a design-mode server. The following example shows the use of the `-dfo` option to delete the `Order` business object in spite of the fact that other components in the system (such as maps and relationships) have references to it:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1
-dfoBusObj:Order
```

*Deleting the referents as well:* Another way you can delete a component that has referents is to use the `-deep` option to delete the referents as well. This deletes the component and all of the components that have references to it. The following example shows the use of the `-deep` option when using the `-do` option to delete the `Customer` business object:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1
-doBusObj:Customer -deep
```

This option, unlike the `-dfo` option, is supported with servers running in production mode because the deletion of the referents along with the component guarantees that the repository remains valid. Keep in mind, however, that it can result in many components being deleted; you should be aware of the implications of this action prior to taking it.

## Exporting components to a package

The `-o` option allows you to export components from the repository to a package. You must specify the name of the package file. When the `-o` option is used alone the entire repository is exported to a file, as in the following example:

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1
-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 10 on page 70, 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 -pnul1
-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_ClarityToSAP` collaboration object and all of its dependencies.

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnul1
-eCollaboration:CustomerSync_ClarityToSAP -oCustomerSyncInterface.jar -deep
```

If you want to export specific components, but do not want to have to enter the entity type keyword and component names, you can store them in a text file and use the `-f` option. This is very convenient when you want to frequently export the same components. The following example uses the `-f` option to load the components listed in a text file named `Components.txt` :

```
C:\WebSphereICS420DEV>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull  
-fComponents.txt -oCustomerSyncInterface.jar -deep
```

The contents of the file `Components.txt` are shown below; a paragraph return follows each entity type keyword and name combination:

```
BusObj:Customer  
Relationship:Customer  
CollabTemplate:CustomerSync
```

**Note:** `Repos_copy` and System Manager are unfortunately inconsistent with respect to what they identify as “dependencies”. If you attempt to delete a component using `repos_copy` but there are components that depend upon it then `repos_copy` lists those referring components as dependencies. However, if you right-click the component in System Manager and select **Show Dependencies** from the context menu the tool lists the components that the selected component depends on.

### Printing a list of components in the repository

You can use the `-summary` argument when executing `repos_copy` to print a list of the components in the repository. The output is presented in XML format. Although it is not particularly useful to view at the command line, you can combine the `-summary` argument with the `-o` argument to redirect the output to a file and then view the file in a browser or XML editor. The command usage in this case would be the following:

```
C:\>repos_copy -sWebSphereICS420DEVServer -uadmin -pnull -summary -oRepository.xml
```

## Locale for `repos_copy` files

The `repos_copy` utility reads metadata from the repository and writes the data out to files in Unicode (UTF-8 format). It also reads such files and loads them into the repository in Unicode (UTF-8 or UCS-2, as the underlying repository database dictates).

`Repos_copy` files created with IBM WebSphere ICS version levels earlier than 4.1.1 can be loaded into the repository correctly only if the dates and times for the component schedules are in full US format. (This is usually not an issue. `Repos_copy` saves all schedule dates in full US format only. The incompatibility could typically arise if the `repos_copy` files have been manually edited.)

---

## Scheduling jobs

Scheduling jobs allows you to create schedules to manipulate the operational states (start, stop, and pause) of connectors and collaborations. By manipulating component states, you can better manage how InterChange Server 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:

“About scheduling jobs” on page 77

“Creating schedules” on page 79

“Modifying schedules” on page 80

“Deleting schedules” on page 81

“Displaying schedules” on page 81

“Enabling or disabling schedules” on page 81

## About scheduling jobs

Scheduling jobs is done through the Schedule window (see Figure 24). 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.

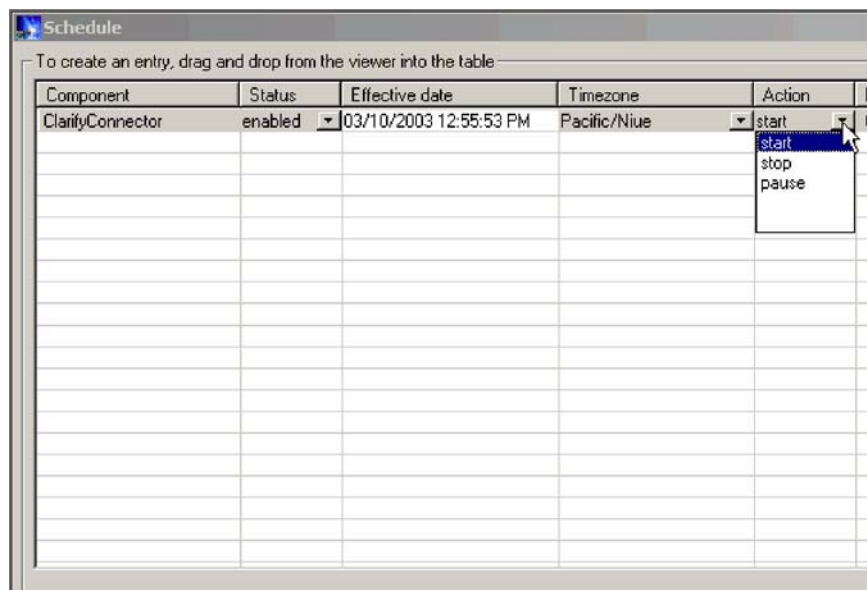


Figure 24. Schedule window

When you create a schedule for a component, you supply information such as when and how often (recurrence) an action (state change) occurs. By default, no schedules are defined for a component. You can define as many schedules as you want for a component. Once a schedule is set, you can enable or disable its use.

The Schedule window allows you to determine the following items:

- |                 |  |
|-----------------|--|
| Status          | Enable turns the schedule on and Disable turns the schedule off. The default status is enabled.  |
| Effective Date  | The date and time the schedule is enabled. The default is the current date and time.   |
| Timezone        | The time zone where the server is located. The default is Pacific Standard time.   |
| Action          | The action the schedule performs. Actions are Start, Pause, and Stop.  |
| Next Occurrence | The next time the scheduled action occurs. If the schedule is non-recurring, the date is the same as the Effective Date. If the schedule is disabled, this field is blank. |



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

If you choose to make the schedule recurring, you can choose from several options including daily, weekly, or monthly.

Because each schedule consists of one action that occurs at a specified time, to create an interval when the server processes a component, you must define both a schedule to start and end processing. As an example, for a connector, you can create one schedule to start processing events at 1 A.M., and another schedule to pause processing at 3 A.M., daily. Only during that two-hour time period can the connector deliver events to InterChange Server Express for processing by collaborations that subscribe to that connector.

### About scheduling connectors

When you schedule the connector operation, the state you select (start, pause, or stop) determines to what extent work is processed. For example, when you start a connector, it constantly polls an application for new events. When you pause a connector, it stops polling until started again, but is still able to handle service call requests from InterChange Server. A stopped connector is inactive.

By manipulating connector activity with collaboration activity, it is possible to schedule dedicated event processing for an application during a specified time window. To do this, both the collaboration and connector must be running during the same time interval. If the connector was paused, events that were queued can be processed when the connector resumes its activity.

### About scheduling collaboration objects

As with connectors, when you schedule the collaboration Object's operation, the state you select (start, pause, or stop) determines to what extent work is processed. To review the collaboration states, see "Collaboration object states" on page 53. 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 25, each collaboration object gets a dedicated time period when the connector is processing only that collaboration object's events.

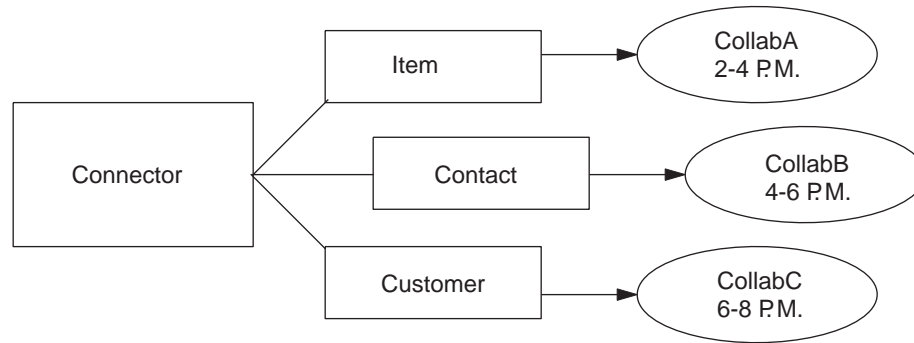


Figure 25. Dedicated processing

Multiple collaboration objects can subscribe to the same business object. In that case, the object is sent to InterChange Server, where it remains until it is picked up by each collaboration object that subscribes to it, when the collaboration object is started after being paused.

### Overriding schedules

Using System Monitor, you can override the state of a scheduled component (for example, start a collaboration object that the scheduler stopped a few minutes ago). Or you can set it to a state to one that the scheduler cannot change. For instance, if a collaboration object is scheduled to pause, you can stop it, not allowing the scheduler to pause it (a collaboration object cannot transition from stop to pause). In such a case, the scheduler does not override the manual change, but logs an error instead.

## Creating schedules

To create a schedule for a collaboration or a connector:

1. Open the Schedule window by right-clicking the Schedule folder in System Manager, then selecting "Edit components' schedule." The Schedule window appears, as shown in Figure 24.

2. From System Manager, select the collaboration object or connector to be scheduled and drag it to the Schedule window.

A new line entry with the name and type of the component is created in the Schedule window (for example, ClarifyConnector (Connector)).

3. Enter information about the schedule by clicking the down arrow in each of the schedule cells:

- a. In the Status field, accept **Enable** to turn the scheduled item on or select **Disable** to turn it off.

An enabled schedule is effective as soon as you click OK or Apply; a disabled schedule is immediately dormant until enabled. When a schedule is disabled, the Next Occurrence cell is blank to indicate there is no scheduled occurrence for this schedule item.

- b. In the Effective Date field, use the calendar to select the date and time when the scheduled item will occur.

By default, the current date and time are set. Use the *MM/DD/YYYY hh:mm:ss* format. A 12 or 24 hour clock is used, based on the Time format configured in the Preferences window, which is available from the Edit menu.

- c. In the Timezone field, select the name of the time zone where the scheduled item is being created, if necessary. By default, the time zone for the scheduled item is set to Pacific Standard Time.  
For example, the schedule for a connector is created in New York (select **Eastern Standard** time) while InterChange Server 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 any comments you may have in the Comments text cell. A maximum of 255 characters is allowed.
4. If you want this schedule to be ongoing, click the Recurrence check box and enter information about the next occurrence of the action. Click one of the radio buttons to determine a style for inputting the recurrence information and use the down arrow menus to select specific date information:
    - The first radio button, Every, specifies a number and a date element, such as every 2 days or every 3 weeks.
    - The second radio button specifies the date in terms of a monthly event by the day of the week, such as the first Tuesday of every month or the fourth Friday of every month.
    - The third radio button specifies the date as the last day of some number of months, such as the last day of every 3 months.

If you do not enable the recurrence option, the Next Occurrence field is blank and the schedule expires after it runs. Consistency checks are made to ensure that only one action is scheduled for a particular component on a given date and time. No checks are performed for scheduling conflicts.

**Note:** InterChange Server Express automatically handles changes between standard and daylight savings time for recurring events.

5. Click either of the Show option check boxes to display specific information about schedules. The Show options are:
  - Show Dependencies, which displays schedules for a Collaboration object's bound connectors and collaborations.
  - Show Expired, which displays schedules that have already processed and whose time to run has expired. Only non-recurring schedules expire.
6. Click OK or Apply to create the scheduled item, which is effective once InterChange Server Express receives the information.

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

**Tip:** To schedule a time interval when events are processed for a component, you must create a schedule with the Start action and another with the Stop or Pause action. See "About scheduling jobs" on page 77 for information and examples about determining start and end schedules.

## Modifying schedules

To modify an existing schedule for a collaboration or a connector:

1. Right-click the component in System Manager, then select "Edit Components' schedule." The Schedule window appears (see Figure 24).

2. Edit any field in the Schedule list window to change its value.  
To edit Recurrence options, click the cursor anywhere on the scheduled item row; the recurrent values for that scheduled item display in the Recurrence pane if they have been assigned.
3. Click OK to save changes and exit, or click Apply to save changes and keep the window open.

## Deleting schedules

To delete an existing schedule for a collaboration object or a connector:

1. Right-click the component from System Manager, then select "Edit Components' schedule." The Schedule window appears (see Figure 24).
2. Select a scheduled item in the schedule list and click the Delete button (or use the keyboard Delete key) to remove the schedule.
3. Click OK to save changes and exit, or click Apply to save changes and keep the window open.

## Displaying schedules

To display a schedule or a group of schedules:

1. Select and open an object for displaying schedules:
  - Collaboration or connector icon from System Manager. The Schedule window displays all schedules defined for that object. If Show Dependencies is active, all schedules for components connected to the object are displayed.
  - Collaboration or connector folder from System Manager. The Schedule window displays schedules for all objects in that folder.
  - InterChange Server Express from System Manager. The Schedule window displays all schedules defined for that server, with times shown for the server's time zone.
  - If you choose the Schedule option from the main window, all schedules in the system are displayed.
2. Click any of the column headings to sort schedules by that column.

## Enabling or disabling schedules

To selectively disable or enable schedules:

1. Select an object for displaying schedules.  
See "Displaying schedules" on page 81.
2. Enable or disable the schedule:
  - To enable or disable all schedules, click either the Enable All or Disable All radio button.
  - To enable or disable a single schedule, click the down arrow in the Status column and choose the Enable or Disable option.
3. Click Apply to complete this task.
4. Click OK to exit.



---

## Notices

IBM may not offer the products, services, or features discussed in this document in all countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing  
IBM Corporation  
North Castle Drive  
Armonk, NY 10504-1785  
U.S.A.

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Burlingame Laboratory Director  
IBM Burlingame Laboratory  
577 Airport Blvd., Suite 800

Burlingame, CA 94010  
U.S.A

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement, or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not necessarily tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

This information may contain examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples may include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

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

---

## Programming interface information

Programming interface information, if provided, is intended to help you create application software using this program.

General-use programming interfaces allow you to write application software that obtain the services of this program's tools.

However, this information may also contain diagnosis, modification, and tuning information. Diagnosis, modification and tuning information is provided to help you debug your application software.

**Warning:** Do not use this diagnosis, modification, and tuning information as a programming interface because it is subject to change.

---

## Trademarks and service marks

The following terms are trademarks or registered trademarks of International Business Machines Corporation in the United States or other countries, or both:

IBM  
the IBM logo  
AIX  
CrossWorlds  
DB2  
DB2 Universal Database  
Domino  
Lotus  
Lotus Notes  
MQIntegrator  
MQSeries  
Tivoli  
WebSphere

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

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

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

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

The Monitor Definition Wizard and System Monitor include software developed by the Eclipse Project (<http://www.eclipse.org>).



IBM WebSphere InterChange Server V4.2.1, IBM WebSphere Business Integration Toolset V4.2.1, IBM WebSphere Business Integration Adapters, V2.3.1, IBM WebSphere Business Integration Collaborations V4.2.









Printed in USA