IBM WebSphere InterChange Server

Problem Determination Guide

V4.3.0

IBM WebSphere InterChange Server

Problem Determination Guide

V4.3.0

Note!

Before using this information and the product it supports, read the information in “Notices” on page 49.

30September2004

This edition of this document applies to IBM WebSphere InterChange Server (5724-178), version 4.3.0, IBM WebSphere Business Integration Toolset (5724-177), version 4.3.0.

To send us your comments about IBM WebSphere Business Integration documentation, e-mail to doc-comments@us.ibm.com. We look forward to hearing from you.

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

Contents

About this document

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

This document contains troubleshooting and problem determination information for the WebSphere InterChange Server system.

Audience

This document is for system administrators, installers, support and service personnel who install and administer the WebSphere InterChange Server product.

Related documents

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

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

- For InterChange Server documentation:

  http://www.ibm.com/websphere/integration/wicserver/infocenter
- For collaboration documentation:

  http://www.ibm.com/websphere/integration/wbicollaborations/infocenter
- For WebSphere Business Integration Adapters documentation:

  http://www.ibm.com/websphere/integration/wbiadapters/infocenter

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

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

Typographic conventions

This document uses the following conventions:

`courier font`	Indicates a literal value, such as a command name, filename, information that you type, or information that the system prints on the screen.
bold	Indicates a new term the first time that it appears.
italic, *`italic`*	Indicates a variable name or a cross-reference.

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

New in this release

New in release 4.3

Version 4.3.0 of the Problem Determination Guide is the first release as part of IBM WebSphere InterChange Server.

Chapter 1. Troubleshooting InterChange Server

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

Using log and trace files for troubleshooting

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

About logging

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

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

Bidirectional (BiDi) characters recorded in log and trace files might be encoded in different code pages which may cause display problems with standard Windows editors. ICS components are configured to dump the trace to STDOUT and if the data is redirected into a file, the encoding of trace messages might be corrupted. If ICS components are run from the DOS prompt, the change codepage command

chcp is required for encoding the traces. For more information on configuring your system for BiDi languages see *System Installation Guide for Windows*.

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

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

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

Note: Log Viewer is an IBM WebSphere ICS system tool, which means it runs only on Windows operating systems. To configure or view a UNIX log file or message using Log Viewer, copy the file from the UNIX machine to a Windows machine that has the IBM WebSphere ICS product installed.

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

For information about viewing logging and tracing messages using Log Viewer, see "Viewing log messages" on page 21.

For background information about tracing, see "About tracing" on page 10.

Collaboration object messages

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

As InterChange Server executes collaboration objects, it writes their messages to its log. For information about configuring system logging, see "Configuring logging and tracing" on page 13.

In addition, you can send a collaboration object's messages by e-mail to one or more recipients. You can specify a separate set of e-mail recipients for each collaboration object. For information on the rules for using e-mail notification, refer to "Configuring e-mail notification of log messages" on page 7.

Connector messages

Connector messages are sent to the InterChange Server message destination. Depending on your operating system, messages appear in one of the following ways:

UNIX

A connector logs messages to STDOUT by default, then those messages are rerouted to `connector_manager_<`*name_of_connector*`>.log`.

Windows

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

For information about connector logging, see “Connector Agent logs” on page 4.

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

Message formats

All messages are formatted so they can easily be filtered. Logged messages for InterChange Server and connectors use the same format, which is described in Table 1. When business objects are configured for flow tracing messages, they use these fields and the additional fields, denoted by an asterisk in Table 1. A message delivered to InterChange Server has the following format, using some or all of the following parameters:

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

Table 1. Message format

Variable	Description
Time	Timestamp: the date of logging in the format *year/month/date time.*
System	The type of component (system identifier). It can be Server, Collaboration, Business Object, or ConnectorAgent.
Thread	Thread name and thread ID
Name	The name of the component, such as ClarifyConnector.
MsgType	Indicates the severity of the message. See Table 2.
MsgID	The message number.
*SubSystem**	The subsystem of the current system. It can be Event Management, Messaging, Repository, or Database Connectivity.
*FIID**	The flow initiator ID of the business object.
*BO**	Business object name.
MsgText	The associated text for the message number.
*BOD**	Business object dump. The data contained in the business object.

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

Table 2 describes the types of IBM WebSphere ICS messages.

Table 2. Message types

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

Table 2. Message types (continued)

Type	Description
Flow Trace	Flow tracing information for business objects.
Internal Error	A serious internal problem that should be investigated.
Internal Fatal Error	An internal error that stops operation. It should be reported.

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

System logs

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

- The system log file, `InterchangeSystem.log`, which you can create in the product directory.
- A log file that you specify.
- Standard output (STDOUT). Depending on your operating system, STDOUT appears in one the following ways:

UNIX

STDOUT messages are redirected to the log file to `$`*`PRODUCTDIR`*`/logs/ics_manager.log`.

Windows

STDOUT messages appear in the Command Prompt window in which InterChange Server starts.

- In addition to managing regular logging, InterChange Server can send an e-mail message to a specified user when it generates error or fatal error messages.

By default, trace messages are sent to the log file. In some cases, this file may become too large, so it is recommended that you specify a separate trace file. See "About tracing" on page 10 for instructions on setting up a separate file for trace messages.

"Configuring logging and tracing" on page 13 describes how to set the destination for logging.

To specify the recipient for e-mail notification, see "Configuring e-mail notification of log messages" on page 7.

"Log/Trace file management" on page 5 describes how to keep log files from becoming too large.

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

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

A connector agent logs messages to a local destination and can also send its messages to InterChange Server for logging. To specify a log file name, edit the `LogFileName` property and insert the name of the log file you want to use. The default log file (located at *ProductDir*/`logs/connector_manager_`*ConnectoName*`.log` directory for UNIX and `STDOUT` for Windows), contains text for the error and informational messages raised from the connector. *Name* is the name of the application. For example, the default message file for the Oracle connector is `OracleConnector.txt`.

Table 3 describes the properties you can edit that determine where a connector agent logs messages.

Table 3. Connector agent log message properties

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

For task instructions on configuring these properties, see "Configuring the connector agent logging destination" on page 17.

High-Availability system logs: The InterChange Server system log for High Availability (HA) is configured to reside on the cluster shared drive (for example, Z:). This log is viewable only from the active system. For information about setting up the log to reside on the shared drive, see the *System Installation Guide for UNIX* or *for Windows*.

The HA system also uses the Application log to provide information about the cluster and its events and failures. Be sure to check the Microsoft MSCS online help for information about this tool.

Log/Trace file management

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

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

`_Arc_` *number*

For example, using 5 archive files, if the log file has the name `InterchangeSystem.log`, the first archive created is named `InterchangeSystem_Arc_01.log`. When the new log file fills up, `InterchangeSystem_Arc_01.log` is renamed as `InterchangeSystem_Arc_02.log`, and the log file is again saved to `InterchangeSystem_Arc_01.log` and so on in a circular fashion, until there are five archive files. If there are five archive files, when a new log file is created, existing archive files are renamed and their numbers incremented so the number of archives matches the number you configured, then the oldest file, whose archive number is 05, is deleted. Figure 1 shows the progression of files using this configuration.

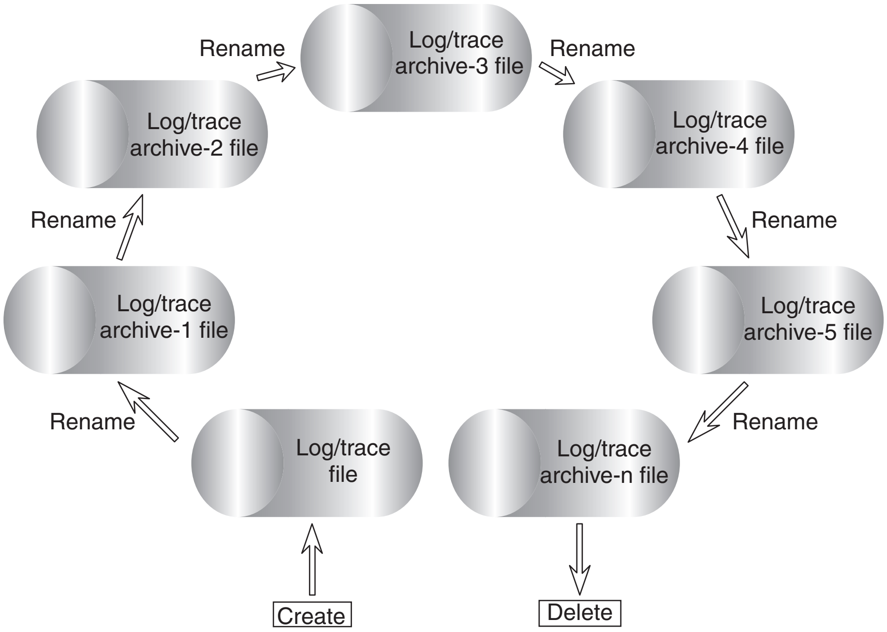

Figure 1. Circular archival logging

See the configuration tasks "Configuring logging and tracing" on page 13 for details.

If the system log file is configured for unlimited size, InterChange Server writes to the log until the disk that the log file is located on gradually fills, and if not administered to, finally produces an error message when the disk is full.

The data in these files should be deleted periodically:

- InterChange System log file. Contains a record of all interactions of the various system components with InterChange Server. The Edit Configuration tool allows you to specify the name and size of this file and a number of archive files for automatically saving versions of this file when it overruns its specified size. See "Configuring InterChange Server logging and tracing destinations" on page 14 for task information.
- Connector log file. Contains flows at the API level that were executed by the connector. Not all applications support these files.

- Connector error log file. Contains error messages returned by the connector. Not all applications support these files.

In addition to the above log files, other log files exist that are specific to each application. Most files are created during runtime if they don't already exist. New information is appended to any existing file. Each component that supplies log information to the files must be taken down before proceeding with a backup.

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

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

Configuring e-mail notification of log messages

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

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

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

- InterChange Server (ICS)
- Collaboration objects
- Connectors

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

Configuring e-mail notification at the system level

Configuring e-mail notification at the system level allows you to configure the server to send e-mail notifications for the following subsystems:

- Default e-mail recipient
- WebSphere MQ errors recipient
- Repository errors recipient
- Database persistence errors recipient
- Database access errors recipient
- Map errors recipient
- Security errors recipient
- Connector controller errors recipient
- DTP errors recipient
- Connector agent errors recipient

To configure e-mail notification at the system level, do the following:

1. From System Manager, right-click the server under Server Instances, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool in which you can edit the `InterchangeSystem.cfg` file.
2. Click the E-mail tab. A dialog box appears in the upper-right section of the System Manager window in which you can enter the parameters necessary for configuring e-mail notification at the system level (see Figure 2).

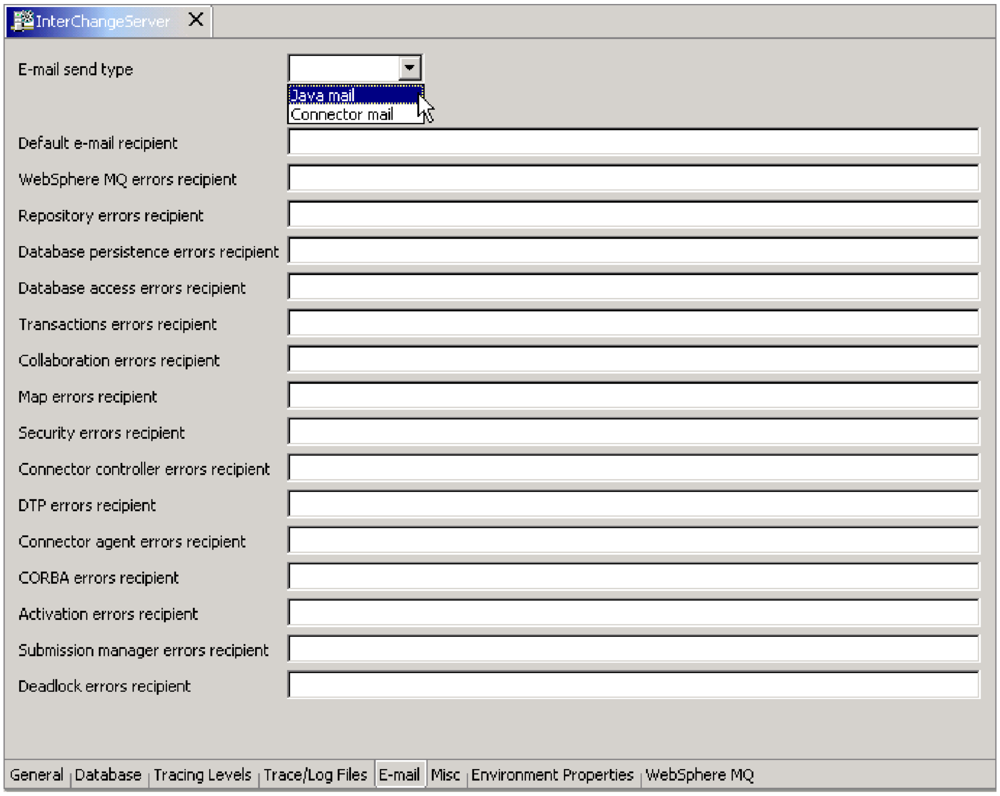
InterChangeServer
E-mail send type
Java mail
Connector mail
Default e-mail recipient
WebSphere MQ errors recipient
Repository errors recipient
Database persistence errors recipient
Database access errors recipient
Transactions errors recipient
Collaboration errors recipient
Map errors recipient
Security errors recipient
Connector controller errors recipient
DTP errors recipient
Connector agent errors recipient
CORBA errors recipient
Activation errors recipient
Submission manager errors recipient
Deadlock errors recipient
General
Database
Tracing Levels
Trace/Log Files
E-mail
Misc
Environment Properties
WebSphere MQ

Figure 2. Edit Configuration, E-mail tab

3. Select an e-mail type from the ″E-mail send type″ drop-down menu. Your choices are ″Java mail″ or ″Connector mail.″
   - Select ″Java mail″ if you do not plan to configure collaborations for e-mail notification. If you plan to configure collaborations for e-mail notification, you must select ″Connector mail.″ If you select ″Java mail,″ an ″SMTP mail host″ field appears at the top of the list of subsystems. Type the host name of the SMTP mail server.

     Note: Be sure to use the correct case, as this field is case sensitive.
   - Select ″Connector mail″ if you plan to configure collaborations for e-mail notification. Selecting ″Connector mail″ requires that you install the e-Mail connector, then configure the e-Mail connector and e-Mail business object. For instructions on installing the e-Mail connector, see the *System Installation Guide for UNIX* or *for Windows*. For instructions on configuring the e-Mail connector and e-Mail business object, see “Configuring the e-Mail business object” on page 9 and “Configuring the e-Mail connector” on page 9.
4. Type one or more valid e-mail addresses in each subsystem field. The address must be SMTP-compliant. For information on SMTP compliancy, see “Using Valid E-mail addresses” on page 10.

5. Select the ″Save <server_name>″ option from the File drop-down menu of System Manager. The system-level e-mail notification information you entered is saved in the `InterchangeSystem.cfg` file.
   - If you chose ″Java mail″ as your e-mail send type, you can skip the following two sections: “Configuring the e-Mail business object” and “Configuring the e-Mail connector.”
   - If you chose ″Connector mail″ as your e-mail send type, you must follow the instructions in the following sections: “Configuring the e-Mail business object” and “Configuring the e-Mail connector.”

Configuring the e-Mail business object: The instructions in this section are required only if you are using the e-Mail connector. Configure the EmailNotification business object to hold the e-mail address of the person who receives e-mail if the intended e-mail recipient is unreachable. As a fail-safe, this should probably be the mail administrator, not the ICS system administrator, to ensure the mail is delivered if the ICS system administrator is unreachable.

1. From System Manager, right-click the EmailNotification business object, then select ″Edit definition.″ The EmailNotification Business Object Designer window appears.
2. From the Attributes tab, type in the return e-mail address in the Defaults column of the `FromAddress` attribute.
3. Select Save from the File drop-down menu.
4. Close Business Object Designer.

Configuring the e-Mail connector: The instructions in this section are required only if you are using the e-Mail connector. To configure the e-Mail connector for sending e-mail:

1. From System Manager, double-click the EmailConnector object. The Connector Configurator window appears.
2. From the Connector Specific Properties tab, double-click the Value cell for the `SMTP_MailHost` property, then type in the name of the Simple Mail Transport Protocol (SMTP) host in the text field.
3. Click Save > To Project from the File drop-down menu.
4. Close the Connector Configurator window.

Configuring e-mail notification at the collaboration object level

This section describes how to configure e-mail notification for a specific collaboration object. Configuring e-mail notification for collaborations requires that you use the e-Mail connector. Be sure to select ″Connector mail″ when configuring e-mail notification at the system level, and be sure to follow the instructions in the following sections to configure the e-Mail connector: “Configuring the e-Mail business object” and “Configuring the e-Mail connector.”

Note: Configuration parameters set at the collaboration object level supersede those set at the system level.

To set an e-mail address to receive messages for a collaboration, do the following:

1. From System Manager, right-click the collaboration object for which you want to configure e-mail notification, then select Properties. The Properties dialog box appears.

2. In the Collaboration General Properties tab, enter a valid e-mail address in the ″Email notification address″ field. The address must be SMTP-compliant. For information on SMTP compliancy, see “Using Valid E-mail addresses.”
3. Click OK to save your changes and close the window.
4. Restart the collaboration for changes to take effect.

Configuring e-mail notification at the connector level

This section describes how to configure e-mail notification for a specific connector.

Note: Configuration parameters set at the collaboration object level supersede those set at the system level.

1. From System Manager, right-click the connector for which you want to configure e-mail notification, then select Edit Definition. The Connector Configurator window appears.
2. From the Properties tab, select `true` in the Value field of the `LogAtInterchangeEnd` property. This enables connector messages to be mailed to the InterChange Server log.
3. Click Save > To Project from the File drop-down menu.
4. Close the Connector Configurator window.
5. Restart the connector for the change to take effect.

Using Valid E-mail addresses

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

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

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

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

About tracing

To troubleshoot a problem, you can turn on tracing. Trace messages help you monitor actions taken in components of the IBM WebSphere ICS system. Trace levels define the amount of detail written to the trace file. The higher the trace level, the more detail you receive. Tracing differs from logging in the following ways:

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

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

Tracing services of InterChange Server

Tracing services for InterChange Server are initially set in parameters of the configuration file for InterChange Server (by default, this file is called `InterchangeSystem.cfg` and resides in the product's top level directory. For details about these parameters, refer to "InterChange Server Configuration Parameters" in the *System Installation Guide for UNIX* or *for Windows*. Settings for these parameters can be updated in the Edit Configuration tool of System Manager, as described in "Configuring tracing levels for InterChange Server, business objects, and collaborations" on page 18.

Tracing collaboration objects

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

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

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

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

For configuration task instructions, see "Configuring tracing levels for InterChange Server, business objects, and collaborations" on page 18.

Tracing connectors

A connector contains two components, the connector controller and the connector agent. The two components can be in different locations on the network and are traced differently. For more information about tracing connectors, see the *Connector Development Guide for* C++ and the *Connector Development Guide for Java*.

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

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

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

4	Prints messages for levels 1 through 3. In addition, traces the passing of business objects between internal levels of the connector.
5	Prints messages for levels 1 through 4. In addition, traces the passing of administrative messages between internal levels of the connector.

A new or changed tracing level takes effect immediately.

For configuration task instructions, see “Configuring connector tracing” on page 19.

Tracing maps

Tracing for IBM WebSphere ICS maps can be done through System Manager. Tracing maps is useful for debugging and keeping track of information as well as error messages created by the map. Tracing of IBM WebSphere ICS maps is turned off by default.

For more information about tracing maps, see the *Map Development Guide.* For configuration task instructions, see “Configuring map tracing” on page 20.

Tracing business objects (flow tracing)

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

Table 4 describes the configurable levels associated with business object tracing:

Table 4. Business object tracing levels

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

For configuration task instructions, see “Configuring tracing levels for InterChange Server, business objects, and collaborations” on page 18.

Tracing Web gateways

Web gateway tracing is provided at two levels, minimal and maximum. This tracing allows you to view information pertaining to whether communication processing is being performed correctly on the gateway. The gateway trace level is set from the Gateway Configuration Tool. For instructions, see the *System Implementation Guide*.

Table 5 on page 13 describes the configurable levels associated with gateway tracing:

Table 5. Gateway tracing levels

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

Tracing SNMP agents

SNMP agent tracing provides trace information for checking the operation of the SNMP agent. The higher the level, the more verbose the output.

Table 6 describes the configurable levels associated with SNMP agent tracing:

Table 6. SNMP agent tracing levels

Level	Description
0	No tracing.
1	Message trace (displays Info, Warning, and Error type messages).
2	Low. Displays all incoming requests and outgoing messages.
3	Medium. Displays information such as Object ID, variable bindings and values.
4	High. Displays commands that access data on ICS and the values that are passed.
5	Method tracing. Errors resulting from ICS interface methods.

Configuring logging and tracing

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

"Opening the Edit Configuration tool" on page 13

"Configuring InterChange Server logging and tracing destinations" on page 14

"Configuring other InterChange Server logging and tracing parameters" on page 16

"Configuring the connector agent logging destination" on page 17

"Configuring tracing levels for InterChange Server, business objects, and collaborations" on page 18

"Configuring the collaboration object trace level" on page 19

"Configuring connector tracing" on page 19

"Configuring map tracing" on page 20

Opening the Edit Configuration tool

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

To open the Edit Configuration tool:

1. From System Manager right-click a server listed under Server Instances, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool from which you can edit the InterchangeServer.cfg file.
2. Click either the Tracing Levels tab or the Trace/Log Files tab to configure tracing and message logging. See "Configuring InterChange Server logging and tracing destinations" and "Configuring tracing levels for InterChange Server, business objects, and collaborations" on page 18 for task instructions.

Configuring InterChange Server logging and tracing destinations

Use the Edit Configuration tool to configure the destination for InterChange Server message logging and tracing.

When configuring these settings, keep the following information about STDOUT in mind:

UNIX

> If you set the logging and tracing to STDOUT, messages are automatically rerouted to `$PRODUCTDIR/logs/ics_manager.log`.

Windows

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

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

   The logging and tracing configuration window appears (see Figure 3).

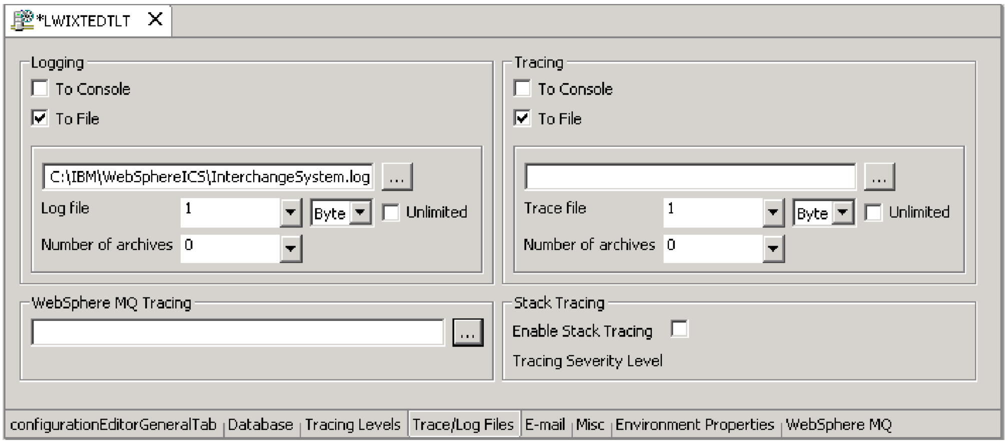

*LWIXTEDTLT

Logging
☐ To Console
☑ To File
C:\IBM\WebSphereICS\InterchangeSystem.log
Log file 1 Byte ☐ Unlimited
Number of archives 0

Tracing
☐ To Console
☑ To File
Trace file 1 Byte ☐ Unlimited
Number of archives 0

WebSphere MQ Tracing

Stack Tracing
Enable Stack Tracing ☐
Tracing Severity Level

Figure 3. Edit Configuration tool, Trace/Log Files tab

2. In the Logging area, select the destination for system logging. You can log to both system console and a log file, but this option should be used only for debugging and on development systems.
   a. To log to the system console (standard output), which is the default, make sure the To console (STDOUT) box is checked. To disable the console as the logging destination, uncheck the To console (STDOUT) option.

b. To log to a file, click the To File box and either type in the full pathname of the file or click the browse (...) button to navigate to the log file.

c. If you select a log file, optionally configure the size of the file in MBs (or keep the default, Unlimited) and the number of archive files to create. For information about archives, see “Maintaining event archives” on page 28.

3. In the Tracing area, select the destination for system tracing.

   a. To log to the system console (standard output), click the To console (STDOUT) box. To disable the console as the logging destination, uncheck the To console (STDOUT) option.

      If you choose to log messages to STDOUT, the messages appear in one of the following ways:

      UNIX

      If you set the logging and tracing to STDOUT, messages are automatically rerouted to `$PRODUCTDIR/logs/ics_manager.log`.

      Windows

      STDOUT appears in the Command Prompt window in which InterChange Server starts.

   b. To log to a file, click the To File box and either type in the full pathname of the file or click the browse (...) button to navigate to the trace file.

   c. If you select a trace file, configure the size of the file in MBs (or keep the default, unlimited checked) and the number of archive files to create. For information about archives, see “Maintaining event archives” on page 28.

4. To configure WebSphere MQ tracing, accept the default pathname of the file or click the browse (...) button to navigate to the log file.

5. Click OK to save changes and exit.

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

Configuring stack trace

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

The following instructions describe how to configure the stack tracing feature.

1. From the InterChage Server Component Management view in System Manager, right-click the ICS instance for which you want to configure stack tracing, then select Edit Configuration.

   The upper-right quadrant of the System Manager becomes a tool in which you can edit the `InterchangeSystem.cfg` file (see Figure 4 on page 16).

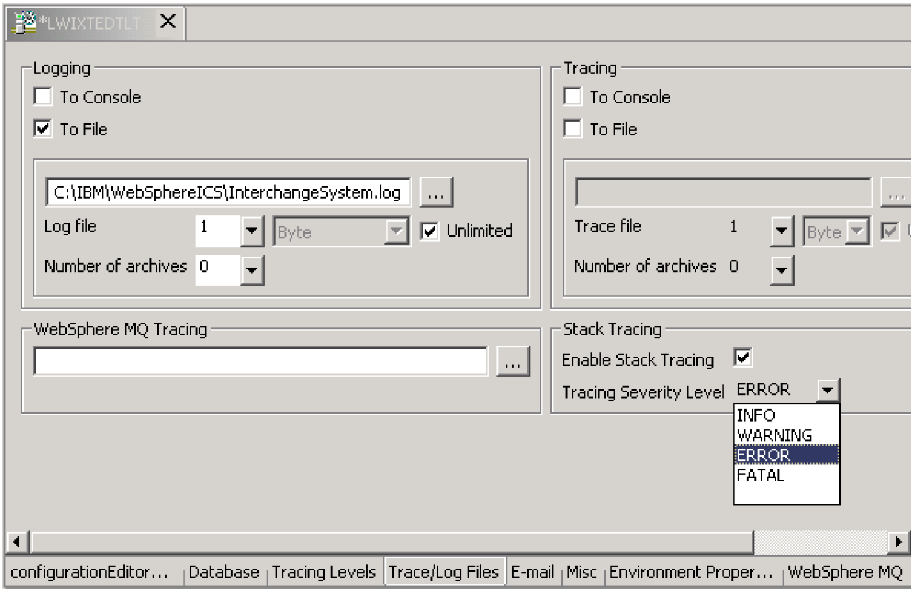
*LWIXTEDTLT

Logging

☐ To Console

☑ To File

C:\IBM\WebSphereICS\InterchangeSystem.log

Log file 1 Byte ☑ Unlimited

Number of archives 0

Tracing

☐ To Console

☐ To File

Trace file 1 Byte ☑ U

Number of archives 0

WebSphere MQ Tracing

Stack Tracing

Enable Stack Tracing ☑

Tracing Severity Level ERROR

Figure 4. Configuring stack tracing

2. In the Trace/Log Files tab, enter the following values:
   - Select the Enable Stack Tracing box. When you enable stack tracing, the Tracing Severity Level option becomes visible.
   - In the Tracing Severity Level drop-down menu, select the severity level for which you want to print the stack tracing. The choices are: `INFO`, `WARNING`, `ERROR`, and `FATAL`. When stack tracing is printed to the log file, it will print the level you select and all levels above it. For example, if you select ERROR, both `ERROR` and `FATAL` messages will be printed to the log file.

Configuring other InterChange Server logging and tracing parameters

You can set other logging and tracing parameters, such as file size and number of archived files. The following instructions describe how to perform these tasks.

To set the maximum size of the log and trace files:

1. From the InterChage Server Component Management view in System Manager, right-click the ICS instance for which you want to set the maximum size for log and trace files, then select Edit Configuration.

   The upper-right quadrant of the System Manager becomes a tool in which you can edit the `InterchangeSystem.cfg` file (see Figure 4).
2. To limit the log file, in the Trace/Log Files tab, select a number to represent the maximum log file size from the Log file drop-down menu, then select Byte, KB (kilobyte), MB (megabyte), or GB (gigabyte).

   Note: The To File and the Unlimited checkboxes in the Logging section must both be selected before you can enter a maximum log file size.
3. To limit the trace file, in the Trace/Log Files tab, select a number to represent the maximum trace file size from the Trace file drop-down menu, then select Byte, KB (kilobyte), MB (megabyte), or GB (gigabyte).

   Note: The To File and the Unlimited checkboxes in the Tracing section must both be selected before you can enter a maximum trace file size.
4. Stop and restart InterChange Server.

To configure the number of archive files:

1. From the InterChage Server Component Management view in System Manager, right-click the ICS instance for which you want to set the maximum size for log and trace files, then select Edit Configuration.

   The upper-right quadrant of the System Manager becomes a tool in which you can edit the `InterchangeSystem.cfg` file (see Figure 4 on page 16).
2. To set the number of log or trace file archives, from the Trace/Log Files tab, select a number from the ″Number of archives″ drop-down list from either the Logging or Tracing section.

   Note: The To File and the Unlimited checkboxes must both be selected before you can select a number from the ″Number of archives″ drop-down menu. This is true for either the Logging or Tracing section.
3. Stop and restart InterChange Server.

Configuring the connector agent logging destination

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

For background information on connector agent logging, see "Connector Agent logs" on page 4.

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

- If you are installing the Remote Agent, you can configure the logging destination during the installation. This option is available for both UNIX and Windows operating systems. For more information about installing Remote Agent, see the *System Implementation Guide*.
- At any other time, you can use Connector Configurator to configure the connector agent logging destination. This can be done only on a Windows machine.

To configure the destination for connector agent logging using Connector Configurator, connect to a server, then follow these steps:

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

1. From System Manager, right-click a connector object, then select Edit Definitions. The Connector Configurator widow appears.
2. From the Standard Properties tab, select one of the logging property values (see Table 3 on page 5 for an explanation of these values), then click Edit.
3. Enter the new value in Value field, then click OK.

   For example, change the `LogAtInterchangeEnd` value to `true` to send messages to the InterChange Server log. If InterChange Server is configured to send e-mail when error and fatal messages are logged, e-mail is sent for the connector agent messages as well.
4. Repeat Steps 3 and 4 to edit the other logging property if necessary.

   For example, to send connector agent messages to a message file instead of the default STDOUT, enter the full pathname of the file in the Value field.
5. Restart the connector for changes to take effect.

Configuring tracing levels for InterChange Server, business objects, and collaborations

This section describes how to configure tracing levels for business objects, collaborations, and the IBM WebSphere ICS subsystems using the Edit Configuration tool in System Manager. For details on viewing the trace messages, see “Log Viewer and tracing” on page 20.

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

To configure tracing:

1. From Edit Configuration tool, select the Tracing Levels tab. Two categories appear in the window: Flow Tracing Levels and IBM WebSphere Business Integration System Tracing Levels (see Figure 5).

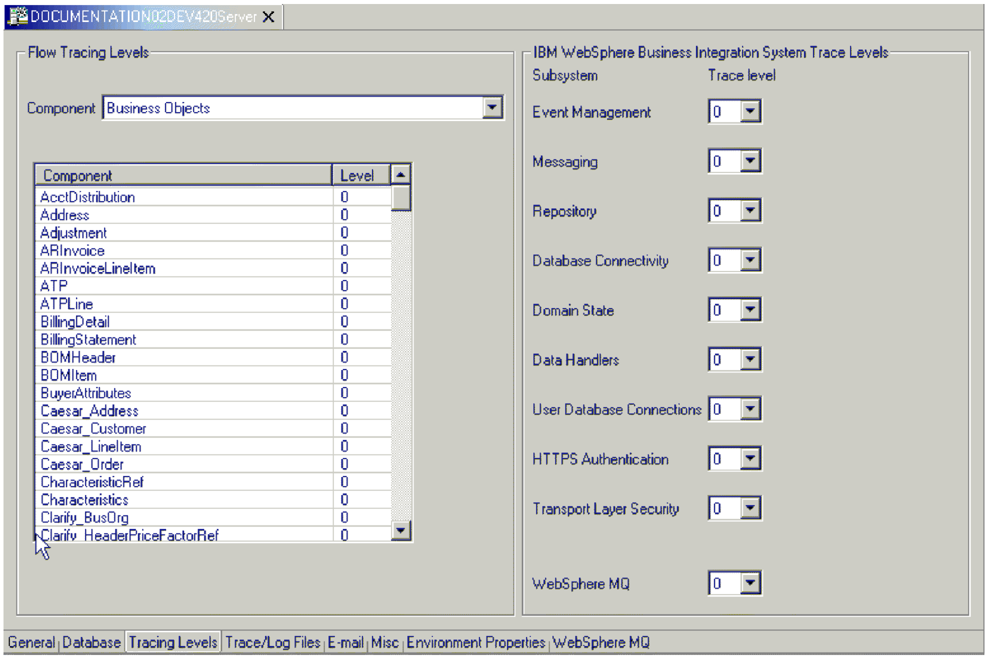

Flow Tracing Levels

Component: Business Objects

Component	Level
AcctDistribution	0
Address	0
Adjustment	0
ARInvoice	0
ARInvoiceLineItem	0
ATP	0
ATPLine	0
BillingDetail	0
BillingStatement	0
BOMHeader	0
BOMItem	0
BuyerAttributes	0
Caesar_Address	0
Caesar_Customer	0
Caesar_LineItem	0
Caesar_Order	0
CharacteristicRef	0
Characteristics	0
Clarify_BusOrg	0
Clarify_HeaderPriceFactorRef	0

IBM WebSphere Business Integration System Trace Levels

Subsystem	Trace level
Event Management	0
Messaging	0
Repository	0
Database Connectivity	0
Domain State	0
Data Handlers	0
User Database Connections	0
HTTPS Authentication	0
Transport Layer Security	0
WebSphere MQ	0

Figure 5. Edit Configuration tool, Tracing Levels tab

2. To configure subsystem tracing, in the IBM WebSphere Business Integration Trace Levels section, click the down arrow next to the subsystems you want to trace and set the trace level using the down arrow menu.

   Setting these trace levels updates parameters in the `InterchangeSystem.cfg` file. For details about what information is produced at the various tracing levels, see “InterChange Server Configuration Parameters” in the *System Installation Guide for UNIX* or *for Windows*.

3. To configure tracing for collaborations, in the Flow Tracing Levels section, click the Component down arrow and select Collaborations.

   The collaboration names that are configured for the system display in the list box.

   Click the down arrow in the Level area to select the trace level for that collaboration.

For a description of the trace levels for collaborations, see "Tracing collaboration objects" on page 11.

4. To configure flow tracing for business objects, click the Component down arrow and select Business Objects. For a description of tracing levels for business objects, see "Tracing business objects (flow tracing)" on page 12.

   The business object names that are configured for the system display in the list box.

   a. Click the down arrow in the Level area to select the trace level for that business object.
5. Click OK to save changes and exit.

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

Configuring the collaboration object trace level

This section describes one of two method that can be used to configure collaboration object trace levels. For the alternative method, see "Configuring tracing levels for InterChange Server, business objects, and collaborations" on page 18.

To configure the runtime system trace level for a collaboration object:

1. From System Manager, right-click the collaboration object name, then select Properties. The Properties dialog box appears.
2. Select a System trace level value, then click OK.

Collaboration object tracing returns messages from inside the collaboration. For a description of collaboration object trace levels, see "Tracing collaboration objects" on page 11.

To configure collaboration object tracing:

1. From System Manager, right-click the collaboration object name, then select Properties. The Properties dialog box appears.
2. In the Collaboration Trace Level field, select a value, then click OK.

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

Configuring connector tracing

For background information, see "Tracing connectors" on page 11.

Connector controller tracing: To trace a connector controller, set the `ControllerTraceLevel` property value to the trace level. Any changes to this property take effect immediately. Trace messages for connector controllers appear wherever InterChange Server sends its trace messages.

To set a connector controller's tracing level, do the following:

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

1. From System Manager, right-click the connector object, then select Edit Definitions. The Connector Configurator window appears.
2. From the Standard Properties tab, click in the Value field of the `ControllerTraceLevel` property, then set the controller tracing level.
3. Select Save > To Project from the File drop-down menu.
4. Close Connector Configurator.

Connector agent tracing: To trace a connector agent, set the `AgentTraceLevel` property value to the trace level. Any changes to this property take effect immediately. Trace messages for connector agents appear wherever the connector agent logs messages. To set a connector agent's tracing level, do the following:

1. From System Manager, right-click the connector, then select Edit Definitions. The Connector Configurator window appears
2. From the Standard Properties tab, click the Value field of the `AgentTraceLevel` property, then set the agent tracing level.
3. Select Save > To Project from the File drop-down menu.
4. Close Connector Configurator.

Configuring map tracing

For background information, see "Tracing maps" on page 12. To set the trace level for a map:

1. From System Manager, right-click the map object, then select Properties. The Maps Property Page appears (see Figure 6).

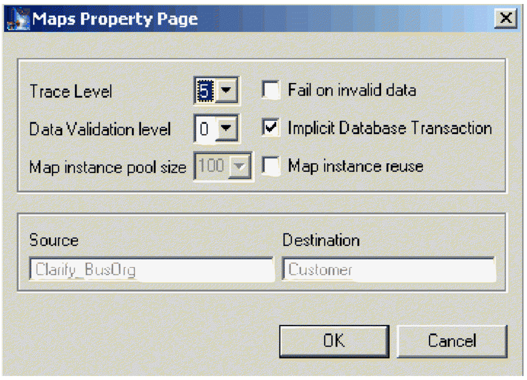

Trace Level 5
Data Validation level 0
Map instance pool size 100
□ Fail on invalid data
☑ Implicit Database Transaction
□ Map instance reuse

Source: Clarify_BusOrg
Destination: Customer

Figure 6. Maps Property Page

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

Using tracing

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

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

For background information about tracing, see "About tracing" on page 10. For configuration information about tracing, see "Configuring logging and tracing" on page 13.

Log Viewer and tracing

Use Log Viewer to display trace information for InterChange Server. In addition to displaying the debugging trace information for collaborations, it allows you to view the progression of a business object as it passes from one processing point to another, for example as the business object exits the collaboration processing point and is sent on to other collaborations or connectors for processing or is forwarded to the mapping stage for data transformation. For information about flow tracing, see "Tracing business objects (flow tracing)" on page 12.

Note: Log Viewer is a WebSphere ICS tool, which means it runs only on Windows machines. To configure or view a UNIX log file using Log Viewer, copy the log file from the UNIX machine to a Windows machine that has the WebSphere ICS product installed.

Viewing log messages

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

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

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

- Use Log Viewer to view system log files in its window, as well as obtain a message explanation interactively by clicking on a message number. See "Using Log Viewer."
- Use a text editor to view a log file and to search the IBM WebSphere ICS message file by message number or text content to obtain an explanation.

Using Log Viewer

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

Note: Log Viewer is a WebSphere ICS tool, which means it runs only on Windows operating systems. To configure or view a UNIX log file using Log Viewer, copy the log file from the UNIX machine to a Windows machine that has the WebSphere ICS product installed.

To start Log Viewer, you can either:

- Select Start > Programs > IBM WebSphere InterChange Server > IBM WebSphere Business Integration Toolset > Administrative > Log Viewer. Use the Open option of the File menu to browse for the log file.
- Use the Run command from the Start menu and browse for the `LogViewer.exe` file. Use the Open option of the File menu to browse for the log file.

Using the Log Viewer menu options, you can perform the following tasks:

- "Setting Log Viewer Preferences" on page 22
- "Changing how messages are viewed" on page 24
- "Manipulating the Log Viewer display output" on page 25
- "Flow tracing a business object" on page 26

Log Viewer, displaying a sample log file is shown in Figure 7 on page 22.

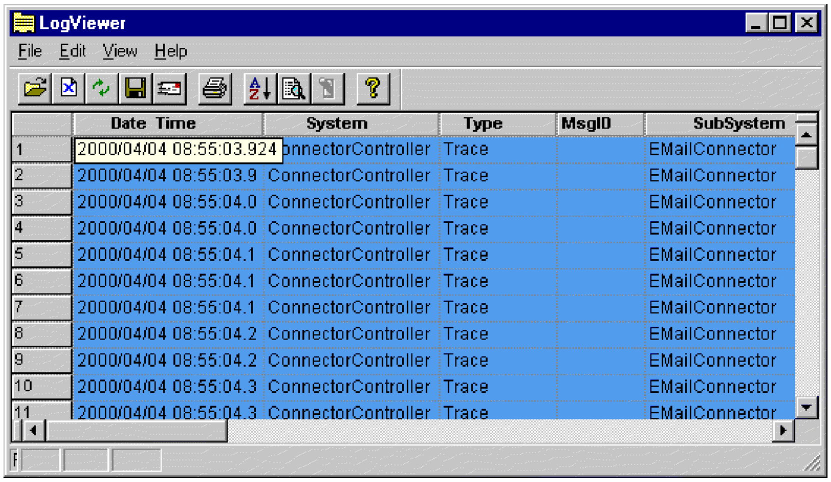

	Date Time	System	Type	MsgID	SubSystem
1	2000/04/04 08:55:03.924	onnectorController	Trace		EMailConnector
2	2000/04/04 08:55:03.9	ConnectorController	Trace		EMailConnector
3	2000/04/04 08:55:04.0	ConnectorController	Trace		EMailConnector
4	2000/04/04 08:55:04.0	ConnectorController	Trace		EMailConnector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

Figure 7. Log Viewer

Setting Log Viewer Preferences:

1. Log Viewer preferences are set by selecting the Preferences option from Log Viewer Edit menu.

   The User Configuration Options, General properties window appears (see Figure 8).

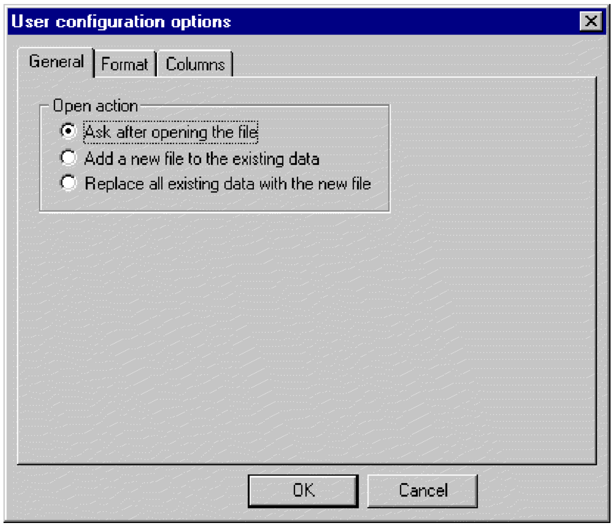

Open action

● Ask after opening the file

○ Add a new file to the existing data

○ Replace all existing data with the new file

Figure 8. Log Viewer User configuration options dialog box, General Properties tab

   This window allows you to determine how to display the log file when you open a log file. The available choices are:

   - Query about what to do each time you open a log file.
   - Merge the log file you are opening with the log file that is currently displayed.
   - Replace the log file that is currently displayed with the contents of the one you are opening.

2. To change the background color and font of the Log Viewer messages, click the Format tab.

The User Configuration Options, Format properties window displays (see Figure 9).

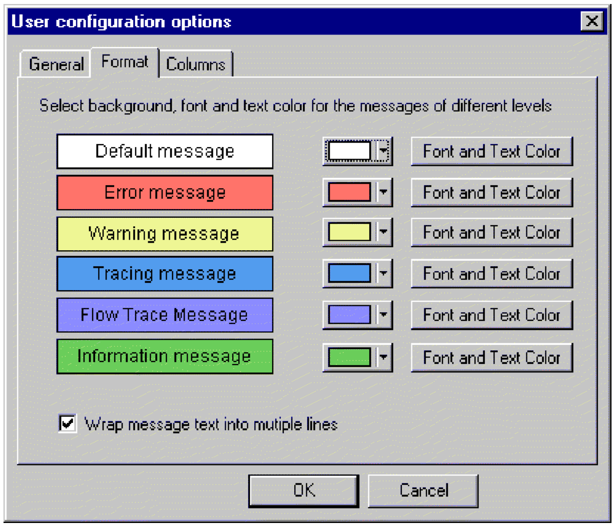

Select background, font and text color for the messages of different levels

☑ Wrap message text into mutiple lines

Figure 9. Log Viewer User configuration options, Format tab

This window allows you to determine how to display the log messages. The available choices are:

- Assign different background colors and fonts for each of the types of messages that display so you can easily recognize their severity (for example, red background with larger font allows for Warning messages).
- Wrap the text of messages if the text is wider than the column.

3. To change which Log Viewer columns display, click the Columns tab.

   The User Configuration Options, Columns properties window displays (see Figure 10).

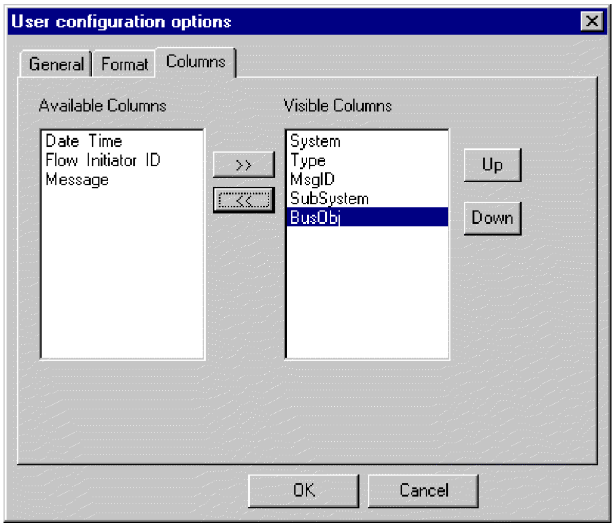

Figure 10. Log Viewer User configuration options, Columns tab

This window allows you to determine which columns display in Log Viewer. To change the columns that display:

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

Changing how messages are viewed: The View menu contains additional options to change Log Viewer displays. In that menu, you'll find the following options:

- Display/hide the Log Viewer toolbar.
- Display/hide the Log Viewer status bar.
- Split the window into two or more views
- Filter or show all messages. Configure filter options by checking filtering options in the filter tabs, such as time range or by type of message (see Figure 11 and Table 2 on page 3). In the Activate Filters area, click the box that is associated with the tab where you selected the filter options, and click OK to enable filtering. The filtered output can be toggled on or off with the Filter Toggle button on the Log Viewer toolbar.

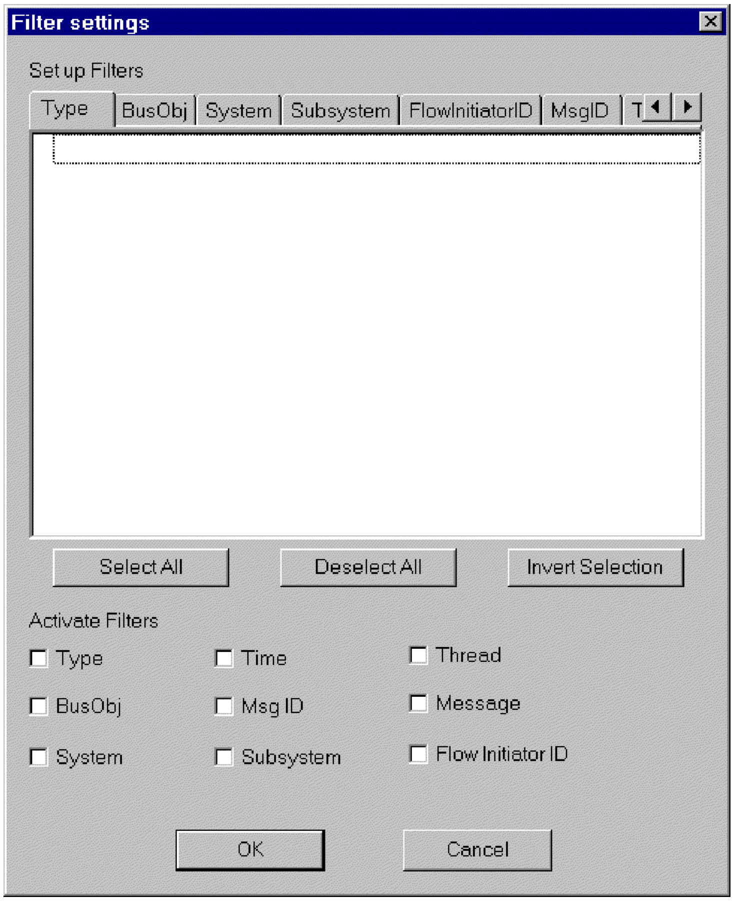

Figure 11. Log Viewer Filter settings Window

- Sort the messages; Figure 12 on page 25 shows the Sort options.

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

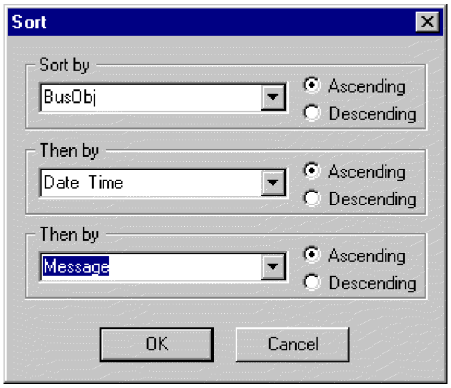

Sort by: BusObj ● Ascending ○ Descending

Then by: Date Time ● Ascending ○ Descending

Then by: Message ● Ascending ○ Descending

Figure 12. Log Viewer Sort Window

Manipulating the Log Viewer display output: Several options are available for manipulating Log Viewer output. In the File menu, there are options for print previewing, printing, saving, refreshing the display, sending to an e-mail recipient, and determining the style for page setup, headers and footers. The variables for header and footers are:

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

Filtering messages: To filter the messages that will be displayed in Log Viewer, choose View->Filter->Use Filter in the Log Viewer menu. The Filter Settings dialog will display.

The Filter Settings dialog displays categories that correspond to the parameters of the logging message format (see Table 2 on page 3 for a descriptive list of the parameters).

In the Filter Settings dialog, you first choose the filtering categories that you want to use, then select the specific items that you want to display from each category, and then choose which filters you want to activate for your current Log Viewer display.

Follow these steps:

1. In the Filter Settings dialog, choose a tab under Set up Filters to display the items that you want to use for filtering messages. For example, choose the BusObj tab to display a list of business objects to be used in filtering, and choose Time if you want to filter according to the timestamp of the message. You can set up multiple filters, and use them either separately or in conjunction with each other.
2. In the displayed list of items, select each item for which you want to view messages in Log Viewer. For example, if you want to view messages related only to the Cost and Customer business objects, select only those business

objects in the list. If you want to view only messages that are timestamped between 5 March 2002 at 9:00 AM and 6 March 2002 at 5:00 PM, select the range for those times under the Time tab.

You can use the buttons below the list box to select all the displayed items, or to deselect all the displayed items, or to invert your current selection choices.

3. Under Activate Filters, check the box for each filter type that you want to activate. For example, if you want to see all messages for the Cost and Customer business objects (which you specified in the previous step), activate only the BusObj filter. If you want to see only those messages for the Cost and Customer business objects that have a particular timestamp, activate both the BusObj filter and the Time filter.
4. Choose OK. The Filter Settings dialog closes, and the Log Viewer display refreshes to show only those messages that you have allowed through the filters.

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

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

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

To trace a business object flow:

1. Set the trace log file destination, if necessary (see "Configuring logging and tracing" on page 13).
2. Select the originating triggering business object (not the generic business object) to trace and set its trace level (see "Configuring tracing levels for InterChange Server, business objects, and collaborations" on page 18).
3. Send some event from the source connector to the destination connector.
4. Open Log Viewer (see "Using Log Viewer" on page 21).
5. Set the display preferences to view the flow tracing (see "Changing how messages are viewed" on page 24).
6. Click any error message button in the MsgID column to view the text of the message.
7. Click any of the business object name buttons in the BusObj column to view the data contained within the business object.

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

Controlling server memory usage

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

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

Perform the following steps to implement memory checking:

- Add the following parameter to the ICS startup script file, as a Java property parameter (-D option):

  ```
  CW_MEMORY_MAX
  ```

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

  ```
  -DCW_MEMORY_MAX=512m
  ```

- Use System Manager to customize server-memory usage by doing the following:
  - In the InterChange Server Component Management view, right click the server instance and click **Edit Configuration**.
  - Click the **Misc** tab and edit the values under **Server Memory**, as show in Figure 13 on page 28:

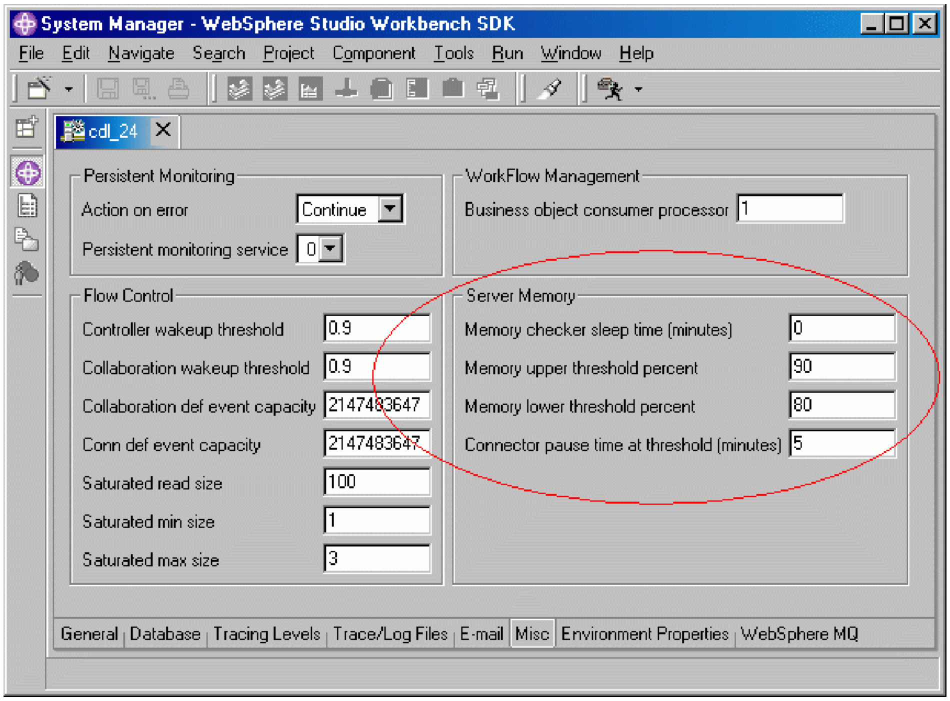
System Manager - WebSphere Studio Workbench SDK
File Edit Navigate Search Project Component Tools Run Window Help
cdl_24
General Database Tracing Levels Trace/Log Files E-mail Misc Environment Properties WebSphere MQ

Persistent Monitoring

Action on error: Continue

Persistent monitoring service: 0

WorkFlow Management

Business object consumer processor: 1

Flow Control

Controller wakeup threshold: 0.9

Collaboration wakeup threshold: 0.9

Collaboration def event capacity: 2147483647

Conn def event capacity: 2147483647

Saturated read size: 100

Saturated min size: 1

Saturated max size: 3

Server Memory

Memory checker sleep time (minutes): 0

Memory upper threshold percent: 90

Memory lower threshold percent: 80

Connector pause time at threshold (minutes): 5

Figure 13. Edit Configuration tool, Misc Tab

Maintaining event archives

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

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

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

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

Managing WIP connections

You should configure enough database Work-in-Progress (WIP) connections to accommodate potential runtime-driven work-load.

To do so, add the following property to the [Event_Management] section of the InterchangeSystem.cfg file:

```
MIN_CONNECTIONS = n
```

Where

```
n=sum_of_all_max_threads_for_collabs_used_simultaneously_plus_num
_connectors
```

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

Setting a sufficient minimum number of WIP connections helps eliminate the possibility of a server crash in circumstances where the server is attempting to reboot, and all connectors are trying to do WIP work while all collaboration threads are undergoing wip-commit.

Managing failures

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

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

"Failure recovery for service calls" on page 29

"Strategies for InterChange Server recovery" on page 30

"Critical errors" on page 32

"Lost connection to application" on page 34

"Unknown connector agent status" on page 34

"Database connection failures" on page 34

"Flow failures" on page 35

Failure recovery for service calls

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

To configure a collaboration to persist a failed service call in-Transit state, set the Persist Service Call In Transit State checkbox of the Collaboration General Properties window.

Attention: If the collaboration is part of a collaboration group, setting this property will implicitly apply the same property setting to all other collaborations in the group. You should not set one collaboration for

Persist Service Call In Transit State unless you want all other collaborations in that collaboration group to have that setting as well.

Attention: For transactional collaborations, it is recommended that you not set this property; leave the Persist Call In Transit State checkbox blank.

Strategies for InterChange Server recovery

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

In release 4.2 of the product, two new features were added that can assist in the efficiency of deferred recovery and asynchronous recovery: flow control and storing business object keys as part of the WIP data. Both features reduce the amount of memory needed during an ICS recovery, and therefore, should decrease significantly the amount of time necessary for ICS to reboot during a recovery.

Storing business object keys as part of the WIP data means that during recovery, the business object key is retrieved without deserializing the business object, thus preventing the need to do an MQ or a database round-trip. Flow control is a service that allows you to configure either system-wide or component-level queue depth parameters in an effort to control the memory demands on ICS.

This section covers the following topics:

"Deferred recovery of collaboration events" on page 30

"Asynchronous recovery" on page 32

Deferred recovery of collaboration events

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

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

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

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

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

The RECOVERY_MODE property has two settings, which have these behaviors when a server failure and reboot occurs:

- Deferred

Events that were in the Working state before the server failure will be transitioned to Deferred. No events for this collaboration will be recovered until you resubmit them manually.

- Always

  Events that were in the Working state before the server failure will be recovered. If there are any existing events that were previously transitioned to the deferred state, they will remain deferred until you resubmit them.

The default setting is Always.

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

To set the Collaboration Recovery mode value when you create a new collaboration object, or to change the Recovery mode value at a later time, you can use the Collaboration General Properties window while InterChange Server is running. To open the window and change the value, do the following:

1. From System Manager, right-click a collaboration object, then choose Properties. The Properties dialog box appears (see Figure 14).
2. In the Collaboration Properties dialog, choose the Collaboration General Properties tab. The following dialog displays:

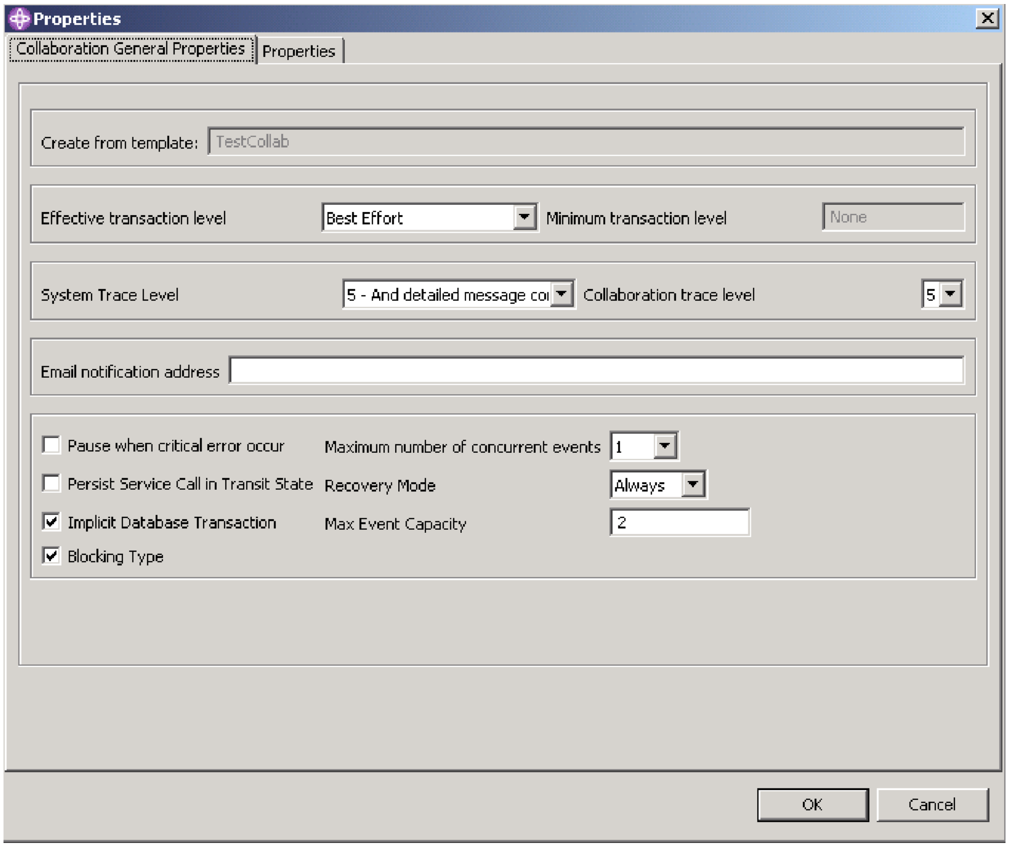

Properties

Collaboration General Properties | Properties

Create from template: TestCollab

Effective transaction level Best Effort

Minimum transaction level None

System Trace Level 5 - And detailed message co

Collaboration trace level 5

Email notification address

☐ Pause when critical error occur

☐ Persist Service Call in Transit State

☑ Implicit Database Transaction

☑ Blocking Type

Maximum number of concurrent events 1

Recovery Mode Always

Max Event Capacity 2

Figure 14. Properties dialog box, Collaboration General Properties tab

3. In the Recovery Mode drop-down menu, choose one of the following:

- Always

  The collaboration will recover all WIP events whose state is WORKING and which it owns at the time of server boot.
- Deferred

  The collaboration will change the WIP event states to a deferred recovery state. You will then need to handle those events at a later time using the Flow Manager. For instructions on using Flow Manager, see *System Administration Guide*.

Asynchronous recovery

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

Critical errors

Critical errors in the WebSphere ICS system can cause problems in your runtime environment. A critical error as defined in the WebSphere ICS system can be generated by one of the following situations:

- Application time-out
- Inability to log on to the application
- Connector agent whose status is unknown

By default, a collaboration continues processing subsequent initiators after a flow has failed. However, a collaboration's behavior can be configured to pause automatically when a critical error occurs that could cause flows to fail. Configuring a collaboration in this way eliminates the possibility of the next flow failing for the same reason by not processing any more initiators after a flow fails. This is critical if the sequence in which initiators are processed needs to be maintained. If the collaboration pauses, the order in which initiators arrived to the server is maintained. At this point, you can fix the critical error, resolve the failed flow, then restart the collaboration. If collaborations are not dependent on an initiator that is associated with a failed flow, you can choose to resume the collaboration and resolve the failed flow at a later time. See "Failed flows" on page 35 for more information on submitting failed events.

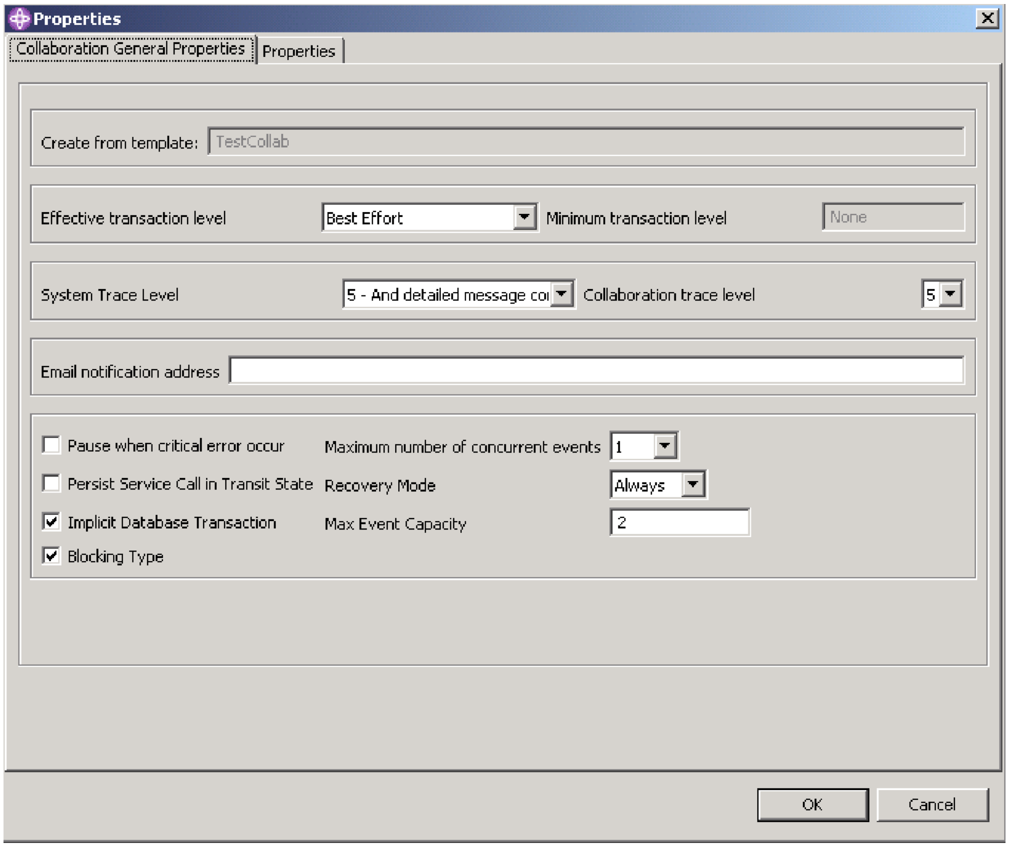

Create from template: TestCollab

Effective transaction level Best Effort
Minimum transaction level None

System Trace Level 5 - And detailed message co
Collaboration trace level 5

Email notification address

☐ Pause when critical error occur
☐ Persist Service Call in Transit State
☑ Implicit Database Transaction
☑ Blocking Type

Maximum number of concurrent events 1
Recovery Mode Always
Max Event Capacity 2

Figure 15. Properties dialog box showing "Pause when critical error occurs" option

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

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

- The connector agent boots up again, automatically notifying the collaboration to resume processing, or
- You manually restart the collaboration, using the Unresolved Flows window

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

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

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

Lost connection to application

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

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

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

Unknown connector agent status

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

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

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

Setting the `ControllerStoreAndForwardMode` property for the connector to `true` makes the connector controller wait for the connector agent to come up before delivering any pending events. Setting this property to `false` makes the connector controller fail collaboration requests. The failed collaboration requests are moved to the resubmission queue and can be resubmitted using Flow Manager. See "Failed flows" on page 35 for more information.

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

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

Database connection failures

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

```
connection in the cache. The maximum number of connections
max-connections-value has been reached.
```

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

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

Flow failures

At times, the WebSphere ICS system or its associated applications may fail. Successfully processing flows that carry data through the WebSphere ICS system is critical, so in a runtime environment it is critical to maintain data consistency. System failures such as system errors, data errors, and critical errors can cause flows to fail to process. The WebSphere ICS system has some built-in capabilities that allow you to handle system failures. The following topics describe two different types of failures:

"Failed flows" on page 35

"Failed transactional collaborations" on page 35

For information on managing, resolving, and preventing flow failures, see *System Administration Guide*.

Failed flows

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

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

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

For instructions on resolving failed flows, see *System Administration Guide*.

Failed transactional collaborations

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

Once a transactional collaboration fails, you need to resolve it. You can handle a failed transactional collaboration by using Flow Manager. For instructions on resolving failed transactional collaborations, see *System Administration Guide*.

Preventing failed transactional collaborations from pausing: The default behavior for a failed transactional collaboration is to pause. You can prevent failed transactional collaborations from pausing by adding a property called `PAUSE_ON_COMPENSATION_FAILURE` to the collaboration template and changing the setting from `TRUE` (default) to `FALSE`. To add the new property and change the setting to `FALSE`, do the following:

1. In System Manager, double-click the collaboration template that failed. Process Designer opens.
2. Double-click the Definitions icon in the left-hand frame. The Template Definitions window appears in the right-hand frame.
3. In the Properties tab, click Add. A Name dialog box appears.
4. Type `PAUSE_ON_COMPENSATION_FAILURE` in the Name field, then click OK. The `PAUSE_ON_COMPENSATION_FAILURE` property appears under the General Properties node in the left-hand pane.
5. With `PAUSE_ON_COMPENSATION_FAILURE` selected in the left-hand pane, select Boolean from the Property Type drop-down menu.
6. Uncheck the IsDefaultVal box for the ″true″ row, and place a check in the IsDefaultVal box in the ″false″ row.
7. Click Apply.
8. Close Process Designer.

Resolving collaboration deadlocks

A deadlock is a situation where two or more processes are unable to proceed because each is waiting for the other processes to proceed. Deadlocks are an undesirable side effect of the concurrency control provided by event isolation within a collaboration. See the *Collaboration Development Guide* for more information on event isolation.

Figure 16 on page 37 illustrates a deadlock between two active collaboration groups resulting from the following sequence of events:

1. At time T1, Collaboration A1 receives an event, E1, then makes a service call to Collaboration B2 and sends a child business object for E1. Collaboration A1 waits for the service call to complete.
2. At time T2, Collaboration B1 receives an event, E2, then makes a service call to Collaboration A2 and sends a child business object for E2. Collaboration A2 waits for the service call to complete.
3. At time T3, Collaboration B2 is waiting for Collaboration B1, since B2 and B1 have the same port binding and the event from B1 was delivered before the event for B2 arrived.
4. At time T4, Collaboration A2 is waiting for Collaboration A1, since A2 and A1 have the same port bindings and the event from A1 was delivered before the event for B1 arrived.

At this point, all collaborations are unable to move forward.

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

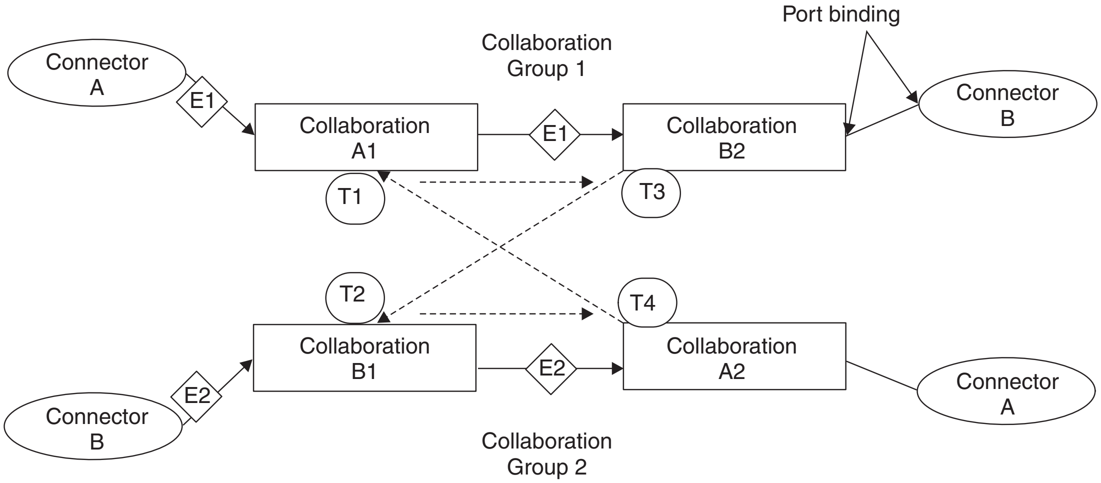

Figure 16. Deadlock between collaboration groups

This section covers the following topics:

"Detecting a collaboration deadlock" on page 37

"Detecting group collaboration deadlocks" on page 38

"Fixing a collaboration deadlock" on page 38

Detecting a collaboration deadlock

You can configure the IBM WebSphere ICS system to either perform or not perform deadlock detection:

- By default, the IBM WebSphere ICS system performs deadlock detection automatically when you start InterChange Server. However, when deadlock detection is performed, startup of InterChange Server can be delayed if a collaboration group contains many collaboration objects, because ICS must traverse all the collaboration objects in the group to determine whether there is a deadlock in the group. This can cause slow startup even when no deadlock exists
- You can configure the IBM WebSphere ICS system to not perform deadlock detection. If you do so, the IBM WebSphere ICS system will start collaboration groups without first checking for deadlocks. This can make it possible for InterChange Server to boot more quickly. However, if deadlock detection is not performed and a deadlock exists, an event that is later sent to a collaboration may fail.

  System Manager does not provide the ability to set the `DEADLOCK_DETECTOR_CHECK` configuration parameter. Instead, to set this configuration parameter, you must edit the `InterchangeSystem.cfg` file and change the parameter's value in this file. Because the `InterchangeSystem.cfg` file is an .xml file, the following lines should exist in this file to define the `DEADLOCK_DETECTOR_CHECK` parameter:

  ```
  <tns:name>DEADLOCK_DETECTOR_CHECK</tns:name>
  <tns:value xml:space="preserve">false</tns:name>
  ```

  To restore deadlock detection, change the `false` value to `true`.

Detecting group collaboration deadlocks

You can check for a group collaboration deadlock in one of the following ways:

- In System Manager, right-click the running group collaboration, then select Diagnostics.

  A window appears with the following message:

  ```
  The following diagnostic tests were run on this collaboration:
  ```

  This message if followed by one of the following results:

  – `Deadlock detection - Fail`

  – `Deadlock detection - Pass`

  – `No diagnostic tests supported at this time`

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

- Check the `InterchangeSystem.log` file for the following error at the time the hanging collaborations were started:

  ```
  Error 11135: Activation of collaboration collaboration_name group could
  cause a potential deadlock with one or more existing collaboration
  groups, and is therefore disallowed.
  ```

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

Fixing a collaboration deadlock

If the WebSphere ICS system encounters a deadlock, you must shut down and restart InterChange Server. First, gracefully shut down all other collaborations, then shut down the server immediately.

Upon system restart, a hung collaboration that caused the deadlock automatically starts and resubscribes to all of the business objects it supports. The business objects that caused the collaborations to enter into a deadlock are redelivered. At this time, the collaborations should not enter into another deadlock because deadlocks are timing dependent. It is unlikely that you will have the exact same server load and isolation sequencing that you had when your system encountered the deadlock.

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

Preventing collaboration deadlocks

You can prevent collaboration deadlocks by configuring the deadlock retry settings in the Database tab of server configuration screen in System Manager. To configure the deadlock retry mechanism, do the following:

1. From System Manager, right-click the server under Server Instances, then select Edit Configuration. The upper-right section of the System Manager window becomes a tool in which you can edit the `InterchangeSystem.cfg` file.
2. Click the Database tab. A dialog box appears in the upper-right section of the System Manager window in which you can enter the parameters necessary for configuring database configuration at the system level (see Figure 17 on page 39).
3. In the ″Max database retry″ field, enter a number that represents the maximum number of retries you want the server to perform if a deadlock occurs.

4. In the "Deadlock retry interval" field, enter a number that represents the number of seconds you want the system to wait before retrying.

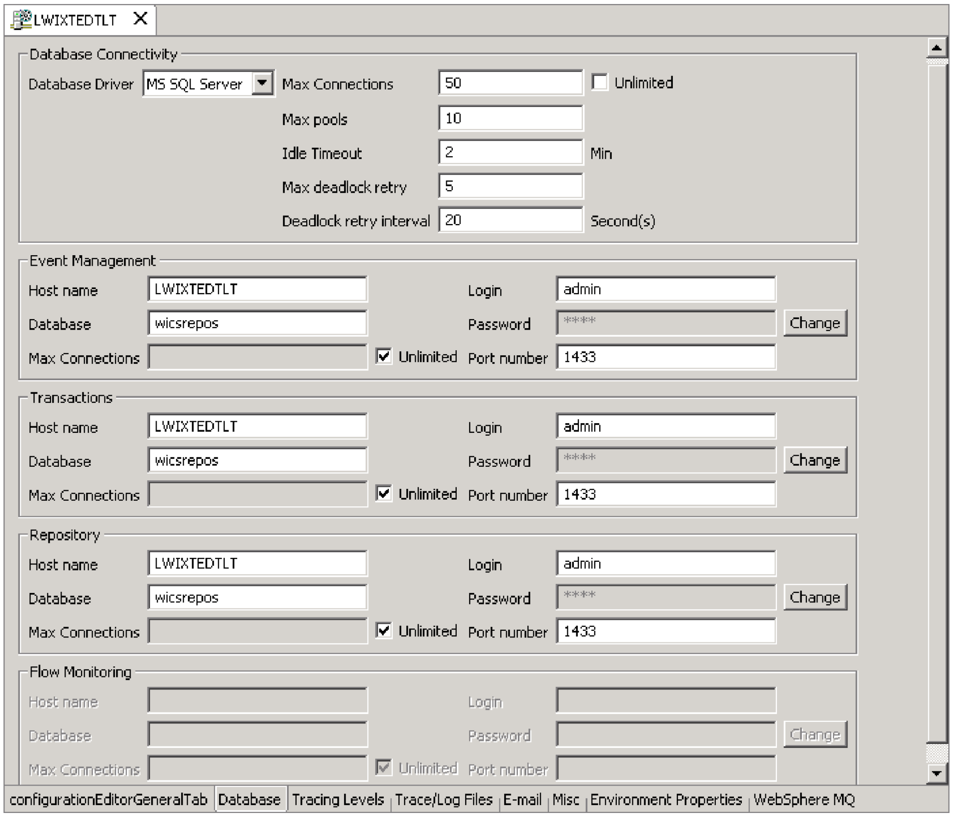

LWIXTEDTLT

Database Connectivity

Database Driver: MS SQL Server
Max Connections: 50 ☐ Unlimited
Max pools: 10
Idle Timeout: 2 Min
Max deadlock retry: 5
Deadlock retry interval: 20 Second(s)

Event Management

Host name: LWIXTEDTLT
Login: admin
Database: wicsrepos
Password: ****
Max Connections: ☑ Unlimited
Port number: 1433

Transactions

Host name: LWIXTEDTLT
Login: admin
Database: wicsrepos
Password: ****
Max Connections: ☑ Unlimited
Port number: 1433

Repository

Host name: LWIXTEDTLT
Login: admin
Database: wicsrepos
Password: ****
Max Connections: ☑ Unlimited
Port number: 1433

Flow Monitoring

Host name:
Login:
Database:
Password:
Max Connections: ☑ Unlimited
Port number:

Figure 17. Edit Configuration screen, Database tab

Managing database connection pools

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

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

If database connection pools have already been set up on your system, you can perform the following tasks to manage how the database connection pool resources are allocated to your collaboration and map processes:

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

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

Chapter 2. Troubleshooting supported software

This chapter provides troubleshooting topics to help determine and resolve problem scenarios that may occur with InterChange Server (ICS) supported software. The following topics are covered:

"Managing WebSphere MQ"

"Display problems when using the Collaboration Events monitor" on page 42

"Troubleshooting WebSphere Studio for Application Developer" on page 43

Managing WebSphere MQ

This section contains information about resolving problems related to WebSphere MQ. For additional information on WebSphere MQ see the InterChange Server installation manuals and the *InterChange Server System Administration Guide.* The following topics are covered:

"Errors related to queue configuration" on page 41

"Failures to create new instances of MQQueueManager" on page 41

"Failure recovery and the queue manager" on page 42

"Memory checker segmentation fault" on page 42

Errors related to queue configuration

WebSphere MQ may report errors if the configuration settings for the queues do not allow for a sufficient number of messages, a sufficient maximum length for messages, or a sufficient size for log files. For information about adjusting these configuration settings, see "Configuring WebSphere MQ Message Queues" in the *System Installation Guide*.

Failures to create new instances of MQQueueManager

On the AIX platform, the use of the bindings transport (C via JNI) in JDK 1.3.1_06 requires some additional configuration to avoid a conflict with shared memory segments. If the configuration is not done, a failure to create a new instance of MQQueueManager, with reason code 2059, can occur. The failure can be accompanied by the generation of FDCs in the `/var/mqm/errors` directory.

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

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

Failure recovery and the queue manager

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

```
AMQ7017 Log not available.
```

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

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

Memory checker segmentation fault

On the AIX platform, the memory checker thread can get a segmentation fault when JMS or other modules are used in WebSphere MQ binding mode.

By default, Java 1.4.2 runs using eight segments. (Eight is the maximum value allowed; each segment is 256 MB). You can reduce the number of segments Java uses with the `LDR_CNTRL` environment variable.

The setting of the LDR_CNTRL=MAXDATA environment variable value is linked to the size of the heap used by the JVM. To simplify the setting of LDR_CNTRL=MAXDATA, the JVM sets an appropriate value that is based on the maximum size of the heap. If LDR_CNTRL=MAXDATA is set before you start the JVM, the JVM uses the specified value. Otherwise, the JVM uses the following algorithm to set LDR_CNTRL=MAXDATA:

On AIX v5.2 or later:

If the heap size is 3 GB or greater, LDR_CNTRL=MAXDATA=0@DSA is set. If the heap size is between 2.3 GB and 3 GB, LDR_CNTR=MAXDATA=0XB0000000@DSA is set. If the heap size is less than 2.3GB, LDR_CNTR=MAXDATA=0XA0000000@DSA is set.

On AIX v5.1:

If the heap size is greater than 1 GB, LDR_CNTRL=MAXDATA is set to an appropriate value. Note that this is an inverse relationship because as the the heap size increases, fewer segments are reserved through the LD_CNTRL=MAXDATA value. For example for a 1 GB heap LDR_CNTRL=MAXDATA is set to 0X60000000, while for a 1.5 GB heap, the LDR_CNTRL=MAXDATA value is 0X40000000. If the heap sizes is smaller than 1 GB, LDR_CNTRL=MAXDATA is set to 0X80000000.

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

Display problems when using the Collaboration Events monitor

This problem occurs for all SVG display types, which are bar, line, stacked bar, and meter.

The problem is that Netscape caches the first SVG chart displayed. The first time you go to one of the views that displays the data in a bar or line, it looks correct. But from then on, it always shows the same chart and data. If you refresh the view or go to another view that has one of the SVG charts (even if it is a different type of chart), it continues to show the same chart and data, because Netscape cached it, and won't refresh even if you reload the page.

The workaround is to disable caching in Netscape so that it loads the SVG charts every time.

Troubleshooting WebSphere Studio for Application Developer

This section describes problems you may encounter when launching WebSphere Studio for Application Developer (WSAD). The following problems are addressed:

- "WSAD will not open"
- "Cannot locate System Manager perspective from WSAD"
- "How to change WSAD workspace directory" on page 44

WSAD will not open

Problem

WSAD will not open.

Solution

One of the initialization of plugins may have failed in a critical way. To verify, look at the log file in `<`*`your WSAS workspace`*`>/.metadata/log`.

Note: The WSAD workspace is typically located at `<`*`your WSAD install directory`*`>/workspace` unless overridded by one of the command line arguments.

If the log file does not indicate any problem, you can launch WSAD with the following command line option:

```
-consoleLog
```

Note 1: The second "`L`" in "consoleLog" must be capitalized.

Note 2: You can supply the command line by editing the shortcut that launches WSAD. This will bring up a console window from which you can watch the initialization process.

Note 3: WSAD is a large application and may take several minutes to launch.

Cannot locate System Manager perspective from WSAD

Problem

WSAD opens, but the System Manager perspective cannot be found.

Solution

The System Manager plugin is probably not loaded properly, or one or more of the components it needs is missing. The following steps describe how to fix this problem:

1. Make sure you have unzipped all of the plugins from the `WebSphereBI\plugins` directory.
2. Look at the WSAD log file, as explained in the "WSAD will not open" section.

How to change WSAD workspace directory

Problem

The WSAD workspace directory needs to be changed.

Solution

Change the following text in the `startcsm.bat` file:

```
-data your_workspace_location
```

Example: `-data C:\myworkspace`

Troubleshooting DB2 server with ICS in multi-threaded mode

This section describes a problem you may encounter when using DB2 database server on Windows when running InterChange Server (ICS) in a multi-threaded mode.

DB2 database server crashes when running ICS in multi-threaded mode

Problem

The DB2 database server instance on Windows crashes when running ICS in multi-threaded mode.

Solution

Check the DB2 compiler options to ensure that they include the correct multi-threaded option: `/MT`. The multi-threaded option may have been incorrectly set to `/mt` for the database compiler. In this case, you may find a message that says that stored procedures have been created in non-fenced mode. Changing this option to `/MT` can correct this problem.

Chapter 3. Troubleshooting InterChange Server tools

This section contains information on problems you may encounter when working with designer tools, System Manager, and other functions or tools of the InterChange Server product.

Troubleshooting designer tools

This section describes problems you may encounter when launching any of the designer tools. These include: Business Object Designer, Map Designer, Process Designer, and Relationship Designer. The following problems are addressed:

- "Designer tool fails to launch and gives a JVM.dll error"
- "Designer tool cannot connect to System Manager"
- "Designer tool cannot be launched from System Manager" on page 46

Designer tool fails to launch and gives a JVM.dll error

Problem

Designer tool fails with the following error message: `failed to locate JVM.dll`.

Solution

The designer tool cannot find the `jvm.dll` file or it is loading an incorrect version of it. One or both of the following solutions should fix the problem:

- Depending on the location of your `jvm.dll` file, make sure you have the correct path listed in your Path System variables. The `jvm.dll` file is typically located in one of the following directories:
  - *ProductDir*`\bin`
  - *ProductDir*`\jre\bin`
  - *ProductDir*`\jre\bin\classic`
- If the above solution does not fix the problem, perform a search on your system for the `jvm.dll` file, and delete all but the one located in *ProductDir*`\jre\bin\classic`.

Designer tool cannot connect to System Manager

Problem

System Manager opens, but designer tools cannot connect to it.

Solution

This typically happens because System Manager cannot bind to a local port (default = 13000). The following list explains three possible reasons why this may happen along with solutions for each explanation.

- There is already another instance of System Manager running. In this case, be sure you are running only one instance of System Manager.
- Another application is using the same port. In this case, you can configure System Manager to use a port other than the default port, which is 13000. To do this, provide the following argument in `startcsm.bat`:

  ```
  -vmargs -Dcom.ibm.btools.internal.comm.port=new_port_number
  ```
- A firewall running on your machine is preventing network connections. In this case, one of the following solutions should allow System Manager to co-exist with your firewall:
  - Change the firewall setting to allow connections to port 13000.
  - Configure System Manager to use a port other than the default port of 13000. Instructions are described in the second bullet-point above.

If the above solutions do not solve the problem, you can try the following debugging argument in the `startcsm.bat` file to see if the console output provides any further insight into the problem:

```
.... -consoleLog -vmargs -Dcom.ibm.btools.internal.comm.debug=true ...
```

Note: If you do not see a separate console window, change the following in `startcsm.bat`:
`-vm %CROSSWORLDS%/jre/bin/javaw.exe`
to
`-vm %CROSSWORLDS%/jre/bin/java.exe`

Designer tool cannot be launched from System Manager

Problem

System Manager opens, but designer tools cannot be launched from it.

Solution

One of the following solutions should fix this problem:

- Make sure your `CROSSWORLDS` system variable is pointing to the correct directory. This variable should be set in the `CWSharedEnv.bat` file, which resides in the `bin` subdirectory of the product directory.
- Make sure that the Path System variable includes a path to the `jvm.dll` file. Example: *ProductDir*`\jre\bin\classic`
- Make sure the `startcsm.bat` file contains the correct value for the `CWTools.home` variable.
  Example: *ProductDir*`\bin`
- Make sure the information contained in the `cwtools.cfg` file is correct. This file is located at *ProductDir*`\bin`.

Troubleshooting problems connecting to InterChange Server in System Manager

There are several common problems that result in System Manager being unable to connect to InterChange Server.

- InterChange Server must be running in order for System Manager to connect to it. Examine the logging output for InterChange Server to ensure that it has a logging statement which reads "*servername*" is ready", where *servername* is the name of your InterChange Server instance.
- The IBM Java Object Request Broker (ORB) must be running for clients such as the tools to communicate with it.
- You must specify the name of InterChange Server exactly as it exists. If you use the wrong case or leave out a single character when trying to register a server instance then System Manager will be unable to connect. You can be sure of the exact name of the server with the following techniques:

  Examine the name exactly as it appears in whatever interface is used to start the server. On Windows, the server name is commonly supplied as a parameter to the start_server.bat batch file. The value is in the Target field of the shortcut that invokes the batch file. On Unix platforms you manually specify the name of the InterChange Server instance.

  Examine the name as it is displayed in the InterChange Server logging output. The log entry "<server name>" is ready" indicates that the server has started and by what name it can be identified. You must specify the correct user name, which is admin.
- You must specify the correct password. The default password is null. If you cannot connect when specifying the password null, make sure that no other developers who work with the server instance have changed the password. If you have cached the user name and password, the password value sometimes gets corrupted. It will still appear as four asterisks, so you will not intuitively think that its value has changed. Delete the cached value in the Password field and type the password again.

No Response from an Eclipse-based tool

If you don't get a response in an Eclipse-based tool (Workbench or System Manager), try restarting the tool.

Tool taking too long to connect to the ICS

If you notice a tool taking a long time to connect to the server, it may be because you have too many users logged into the server. If this happens your administrator can log off idle users and then try again to connect the tool. Or you could have your users logoff and close the tools when they are no longer using them.

Notices

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

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

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

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

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

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

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

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

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

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

Burlingame, CA 94010
U.S.A

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

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

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

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

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

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

COPYRIGHT LICENSE

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

Programming interface information

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

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

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

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

Trademarks and service marks

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

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

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

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

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

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

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

IBM WebSphere InterChange Server version 4.3.0, IBM WebSphere Business Integration Toolset version 4.3.0.

Index

Printed in USA