

WebSphere Process Server for z/OS



Troubleshooting and Support

Version 7.0.0

30 April 2010

This edition applies to version 7, release 0, modification 0 of WebSphere Process Server for z/OS (product number 5655-N53) and to all subsequent releases and modifications until otherwise indicated in new editions.

To send us your comments about this document, send an e-mail message 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.

© **Copyright IBM Corporation 2006, 2010.**

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

Contents

Overview of troubleshooting	1	Finding failed events	38
Troubleshooting checklist for WebSphere Process Server	5	Working with data in failed events	40
Messages overview	7	Resubmitting failed events	45
WebSphere Process Server log files	9	Managing failed SCA events.	47
Transaction log file	11	Managing failed JMS events	48
Troubleshooting the installation	13	Managing failed WebSphere MQ events	49
Message reference for WebSphere Process Server for z/OS installation and configuration	13	Managing stopped Business Process Choreographer events	51
WebSphere Process Server log files	13	Finding business process instances related to a failed event	52
Troubleshooting a failed deployment	15	Finding Common Base Events related to a failed event	52
Deleting JCA activation specifications.	16	Deleting failed events	53
Deleting SIBus destinations	17	Troubleshooting the failed event manager	53
Troubleshooting administration tasks and tools.	19	Recovering from a failure.	55
Profile-specific log files	19	Overview of the recovery process	55
Troubleshooting the failed event manager	21	Triggers for recovery	55
Troubleshooting store-and-forward processing.	22	Assessing the state of the system	56
Troubleshooting the business rules manager	24	Recovery: Analyzing the problem	58
Resolving login errors	24	Situational analysis	59
Resolving login conflict errors	25	Recovery: First steps	59
Resolving access conflict errors	25	Failed-event locations: Where does the data go?	60
Troubleshooting WebSphere Application Server.	27	Use case: recovering data from failed events	61
Tools for troubleshooting your applications	29	Recovery troubleshooting tips	69
Debugging applications in WebSphere Integration Developer	29	Restarting deployment environments	69
Using logging, tracing, and monitoring in your applications	29	Viewing the service integration bus	70
Troubleshooting Service Component Architecture processing and call chains	30	Capturing javacore	74
Managing failed events	33	Servers and recovery mode processing	75
Security considerations for recovery	37	Retention queues and hold queues	76
		Business Process Choreographer maintenance and recovery scripts	77
		Resolving indoubt transactions	79
		Reviewing DB2 diagnostic information	81
		Process recovery troubleshooting tips.	82
		About recovering the messaging subsystem	82
		IBM Support Assistant.	83
		Searching knowledge bases	85
		IBM Support Assistant	87
		Getting fixes	89
		Contacting IBM Software Support	91

Overview of troubleshooting

Troubleshooting is a systematic approach to solving a problem. The goal is to determine why something does not work as expected and how to resolve the problem.

The first step in the troubleshooting process is to describe the problem completely. Without a problem description, neither you or IBM® can know where to start to find the cause of the problem. This step includes asking yourself basic questions, such as:

- What are the symptoms of the problem?
- Where does the problem occur?
- When does the problem occur?
- Under which conditions does the problem occur?
- Can the problem be reproduced?

The answers to these questions typically lead to a good description of the problem, and that is the best way to start down the path of problem resolution.

What are the symptoms of the problem?

When starting to describe a problem, the most obvious question is "What is the problem?" Which might seem like a straightforward question; however, you can break it down into several more-focused questions that create a more descriptive picture of the problem. These questions can include:

- Who, or what, is reporting the problem?
- What are the error codes and messages?
- How does the system fail? For example, is it a loop, hang, lock up, performance degradation, or incorrect result?
- What is the business impact of the problem?

Where does the problem occur?

Determining where the problem originates is not always simple, but it is one of the most important steps in resolving a problem. Many layers of technology can exist between the reporting and failing components. Networks, disks, and drivers are only a few components to be considered when you are investigating problems.

The following questions can help you to focus on where the problem occurs in order to isolate the problem layer.

- Is the problem specific to one platform or operating system, or is it common for multiple platforms or operating systems?
- Is the current environment and configuration supported?

Remember that if one layer reports the problem, the problem does not necessarily originate in that layer. Part of identifying where a problem originates is understanding the environment in which it exists. Take some time to completely describe the problem environment, including the operating system and version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported configuration; many

problems can be traced back to incompatible levels of software that are not intended to run together or have not been fully tested together.

When does the problem occur?

Develop a detailed timeline of events leading up to a failure, especially for those cases that are one-time occurrences. You can most simply do this by working backward: Start at the time an error was reported (as precisely as possible, even down to the millisecond), and work backward through the available logs and information. Typically, you need to look only as far as the first suspicious event that you find in a diagnostic log; however, this is not always simple to do and takes practice. Knowing when to stop looking is especially difficult when multiple layers of technology are involved, and when each has its own diagnostic information.

To develop a detailed timeline of events, answer the following questions:

- Does the problem happen only at a certain time of day or night?
- How often does the problem happen?
- What sequence of events leads up to the time that the problem is reported?
- Does the problem happen after an environment change, such as upgrading or installing software or hardware?

Responding to these types of questions can provide you with a frame of reference in which to investigate the problem.

Under which conditions does the problem occur?

Knowing what other systems and applications are running at the time that a problem occurs is an important part of troubleshooting. These and other questions about your environment can help you to identify the root cause of the problem:

- Does the problem always occur when the same task is being performed?
- Does a certain sequence of events need to occur for the problem to surface?
- Do any other applications fail at the same time?

Answering these types of questions can help you explain the environment in which the problem occurs and correlate any dependencies. Remember that just because multiple problems might have occurred around the same time, the problems are not necessarily related.

Can the problem be reproduced?

From a troubleshooting standpoint, the "ideal" problem is one that can be reproduced. Typically with problems that can be reproduced, you have a larger set of tools or procedures at your disposal to help you investigate. Consequently, problems that you can reproduce are often simpler to debug and solve. However, problems that you can reproduce can have a disadvantage: If the problem is of significant business impact, you do not want it to recur! If possible, re-create the problem in a test or development environment, which typically offers you more flexibility and control during your investigation.

Tip: Simplify the scenario to isolate the problem to a suspected component.

The following questions can help you with reproducing the problem:

- Can the problem be re-created on a test machine?

- Are multiple users or applications encountering the same type of problem?
- Can the problem be re-created by running a single command, a set of commands, a particular application, or a stand-alone application?

Troubleshooting checklist for WebSphere Process Server

Asking questions about hardware and software requirements, product fixes, specific problems, error messages, and diagnostic data can help you troubleshoot WebSphere® Process Server.

The following questions can help you to identify the source of a problem that is occurring with WebSphere Process Server:

1. Is the configuration supported?

Refer to the requirements for WebSphere Process Server to ensure that your system meets all hardware, operating system, and software requirements: [WebSphere Process Server system requirements Web site](#).

2. Have you applied the latest fixes?

3. What is the problem?

- Installing and configuring WebSphere Process Server
- Migrating existing applications and configuration information to WebSphere Process Server
- Deploying applications on WebSphere Process Server
- Administering applications and components on WebSphere Process Server
- Using WebSphere Application Server capabilities in WebSphere Process Server

4. Have any error messages been issued?

5. For additional help in finding error and warning messages, interpreting messages, and configuring log files, see [Diagnosing problems with message logs](#) in the [WebSphere Application Server information center](#).

6. Difficult problems can require the use of tracing, which exposes the low-level flow of control and interactions between components. For help in understanding and using traces, see [Working with trace](#) in the [WebSphere Application Server information center](#).

7. If the checklist does not guide you to a resolution, you can collect additional diagnostic data. This data is necessary for IBM Support to effectively troubleshoot and assist you in resolving the problem. For more information, see ["Contacting IBM Software Support"](#) on page 91.

Messages overview

When you receive a message from WebSphere Process Server, you can often resolve the problem by reading the entire message text and the recovery actions that are associated with the message.

You can find the full text of runtime messages, their explanations, and the recommended recovery actions by searching for the message identifier in the Messages section of the WebSphere Process Server Reference documentation.

Messages displayed during WebSphere Process Server product installation and profile creation are documented in topics under Related Topics at the bottom of this page..

Runtime message identifiers consist of a four- or five-character message prefix, followed by a four- or five-character message number, followed by a single-letter message type code. For example, zzzzL1042C. The message type code describes the severity of the error message, as follows:

- C** Indicates a severe message.
- E** Indicates an urgent message.
- I** Indicates an informational message.
- N** Indicates an error message.
- W** Indicates a warning message.

Related reference

“Message reference for WebSphere Process Server for z/OS installation and configuration” on page 13

The message reference for WebSphere Process Server for z/OS[®] lists the message codes that display while running the install script or when running the configuration script.

WebSphere Process Server log files

There are two distinct groups of log files in the installed product. Logs detailing the product installation, product updates, and profile management are one group. Logs detailing the characteristics and runtime activities of individual profiles make up the second group.

Various log files are created during the installation and uninstallation of WebSphere Process Server and during profile creation, augmentation, and deletion. Examine these logs when problems occur during the product installation and configuration process. The log files and their locations within the product installation are detailed in the "Installation and profile creation log files" topic.

There are also a number of log files that are created for each profile. Some of these logs describe the parameters used for the creation of the profile. These types of log files generally remain unchanged when the profile is fully configured. Other profile-specific logs are continually updated to capture error, warning, and information messages emitted during runtime. Some of these log files are also used to capture a Common Base Event (that may include business object data) that is selected for monitoring. This set of logs is described in "Profile-specific log files" topic.

Related concepts

"Profile-specific log files" on page 19

There are log files detailing the characteristics and runtime activities of individual profiles. These log files are located within the logs directory of the profile path.

Transaction log file

The transaction (tranlog) log file stores critical transactional data that is written to databases. It is an internal file that WebSphere Application Server uses to manage in-flight transactions and attempt to recover them should the server lock up.

DO NOT delete the transaction log file from a production environment. Deleting this file removes information on in-flight transactions from WebSphere Process Server memory. Without the transaction log file, there is no functionality to recover transactional information. In addition, long-running processes remain in an inconsistent state and you cannot complete the process flow except by deleting running instances. Deleting running instances might cause you to lose operational or business-critical data, which makes the database inconsistent with the message destination. Other inconsistencies that may be caused by deleting the transaction log file includes the following:

- Started transactions will neither be rolled back nor committed
- Artifacts will remain in the Java™ Virtual Machine (JVM) since they are referenced or allocated by a transaction but never garbage collected
- Database content (amongst others navigation state of long running BPEL processes) remains in the Business Process Choreographer related tables and are never deleted
- Navigation messages of the Business Process Engine (BPE) of long running processes are never processed further
- Service Component Architecture (SCA) messages that belong to a process navigation and transaction remain on SCA-related queues

Note: Deleting the transaction log from a development environment causes the same problems. Because you can re-create business processes, deleting the files from a test environment is not as damaging as deleting them from a production environment.

Troubleshooting the installation

You can diagnose problems when the installation of WebSphere Process Server is unsuccessful.

Message reference for WebSphere Process Server for z/OS installation and configuration

The message reference for WebSphere Process Server for z/OS lists the message codes that display while running the install script or when running the configuration script.

About the installation error messages

Use the data in the Explanation and user response fields to troubleshoot the WebSphere Process Server for z/OS message codes.

The message code displays as CWPIZyyyyz, where:

- CWPIZ = The WebSphere Process Server for z/OS message prefix
- yyyy = The numeric identifier assigned to the number
- z = Descriptor (E, I or W) for the type of message, where:
 - E = Error message
 - I = Informational message
 - W = Warning message

For a listing of the WebSphere Process Server for z/OS installation error messages, see CWPIZ in the Messages portion of the Reference documentation.

The WebSphere Process Server for z/OS installation error messages are written to the zSMPInstall.log file in the run-time directory. The standard default location for the log file is /WebSphere/V7R0/AppServer/logs/wbi/zSMPInstall.log.

The WebSphere Process Server for z/OS configuration error messages are written to the zWPSConfig.log file and the zWESBConfig.log file in the run-time directory. The standard default location for these log files are /WebSphere/V7R0/AppServer/logs/wbi/zWESBConfig.log and /WebSphere/V7R0/AppServer/logs/wbi/zWPSConfig.log respectively.

Related concepts

“Messages overview” on page 7

When you receive a message from WebSphere Process Server, you can often resolve the problem by reading the entire message text and the recovery actions that are associated with the message.

WebSphere Process Server log files

There are two distinct groups of log files in the installed product. Logs detailing the product installation, product updates, and profile management are one group. Logs detailing the characteristics and runtime activities of individual profiles make up the second group.

Various log files are created during the installation and uninstallation of WebSphere Process Server and during profile creation, augmentation, and deletion. Examine these logs when problems occur during the product installation and configuration process. The log files and their locations within the product installation are detailed in the "Installation and profile creation log files" topic.

There are also a number of log files that are created for each profile. Some of these logs describe the parameters used for the creation of the profile. These types of log files generally remain unchanged when the profile is fully configured. Other profile-specific logs are continually updated to capture error, warning, and information messages emitted during runtime. Some of these log files are also used to capture a Common Base Event (that may include business object data) that is selected for monitoring. This set of logs is described in "Profile-specific log files" topic.

Related concepts

"Profile-specific log files" on page 19

There are log files detailing the characteristics and runtime activities of individual profiles. These log files are located within the logs directory of the profile path.

Troubleshooting a failed deployment

This topic describes the steps to take to determine the cause of a problem when deploying an application. It also presents some possible solutions.

Before you begin

This topic assumes the following things:

- You have a basic understanding of debugging a module.
- Logging and tracing is active while the module is being deployed.

About this task

The task of troubleshooting a deployment begins after you receive notification of an error. There are various symptoms of a failed deployment that you have to inspect before taking action.

Procedure

1. Determine if the application installation failed.

Examine the `SystemOut.log` file for messages that specify the cause of failure. Some of the reasons an application might not install include the following:

- You are attempting to install an application on multiple servers in the same Network Deployment cell.
- An application has the same name as an existing module on the Network Deployment cell to which you are installing the application.
- You are attempting to deploy Java EE modules within an EAR file to different target servers.

Important: If the installation has failed and the application contains services, you must remove any SIBus destinations or JCA activation specifications created before the failure before attempting to reinstall the application. The simplest way to remove these artifacts is to click **Save > Discard all** after the failure. If you inadvertently save the changes, you must manually remove the SIBus destinations and JCA activation specifications (see [Deleting SIBus destinations](#) and [Deleting JCA activation specifications](#) in the [Administering](#) section).

2. If the application is installed correctly, examine it to determine if it started successfully.

If the application did not start successfully, the failure occurred when the server attempted to initiate the resources for the application.

- a. Examine the `SystemOut.log` file for messages that will direct you on how to proceed.
- b. Determine if resources required by the application are available and/or have started successfully.

Resources that are not started prevent an application from running. This protects against lost information. The reasons for a resource not starting include:

- Bindings are specified incorrectly
- Resources are not configured correctly
- Resources are not included in the resource archive (RAR) file

- Web resources not included in the Web services archive (WAR) file
- c. Determine if any components are missing.

The reason for missing a component is an incorrectly built enterprise archive (EAR) file. Make sure that the all of the components required by the module are in the correct folders on the test system on which you built the Java archive (JAR) file. “Preparing to deploy to a server” contains additional information.
 3. Examine the application to see if there is information flowing through it.

Even a running application can fail to process information. Reasons for this are similar to those mentioned in step 2b on page 15.

 - a. Determine if the application uses any services contained in another application. Make sure that the other application is installed and has started successfully.
 - b. Determine if the import and export bindings for devices contained in other applications used by the failing application are configured correctly. Use the administrative console to examine and correct the bindings.
 4. Correct the problem and restart the application.

Deleting JCA activation specifications

The system builds JCA application specifications when installing an application that contains services. There are occasions when you must delete these specifications before reinstalling the application.

Before you begin

If you are deleting the specification because of a failed application installation, make sure the module in the Java Naming and Directory Interface (JNDI) name matches the name of the module that failed to install. The second part of the JNDI name is the name of the module that implemented the destination. For example in `sca/SimpleBOCrsmA/ActivationSpec`, **SimpleBOCrsmA** is the module name.

Required security role for this task: When security and role-based authorization are enabled, you must be logged in as administrator or configurator to perform this task.

About this task

Delete JCA activation specifications when you inadvertently saved a configuration after installing an application that contains services and do not require the specifications.

Procedure

1. Locate the activation specification to delete.

The specifications are contained in the resource adapter panel. Navigate to this panel by clicking **Resources > Resource adapters**.

 - a. Locate the **Platform Messaging Component SPI Resource Adapter**.

To locate this adapter, you must be at the **node** scope for a standalone server or at the **server** scope in a deployment environment.
2. Display the JCA activation specifications associated with the Platform Messaging Component SPI Resource Adapter.

Click on the resource adapter name and the next panel displays the associated specifications.

3. Delete all of the specifications with a **JNDI Name** that matches the module name that you are deleting.
 - a. Click the check box next to the appropriate specifications.
 - b. Click **Delete**.

Results

The system removes selected specifications from the display.

What to do next

Save the changes.

Deleting SIBus destinations

Service integration bus (SIBus) destinations are used to hold messages being processed by SCA modules. If a problem occurs, you might have to remove bus destinations to resolve the problem.

Before you begin

If you are deleting the destination because of a failed application installation, make sure the module in the destination name matches the name of the module that failed to install. The second part of the destination is the name of the module that implemented the destination. For example in `sca/SimpleBOCrsmA/component/test/sca/cros/simple/cust/Customer`, **SimpleBOCrsmA** is the module name.

Required security role for this task: When security and role-based authorization are enabled, you must be logged in as administrator or configurator to perform this task.

About this task

Delete SIBus destinations when you inadvertently saved a configuration after installing an application that contains services or you no longer need the destinations.

Note: This task deletes the destination from the SCA system bus only. You must remove the entries from the application bus also before reinstalling an application that contains services (see Deleting JCA activation specifications in the Administering section of this information center).

Procedure

1. Log into the administrative console.
2. Display the destinations on the SCA system bus.
 - a. In the navigation pane, click **Service integration** → **buses**
 - b. In the content pane, click **SCA.SYSTEM.cell_name.Bus**
 - c. Under Destination resources, click **Destinations**
3. Select the check box next to each destination with a module name that matches the module that you are removing.
4. Click **Delete**.

Results

The panel displays only the remaining destinations.

What to do next

Delete the JCA activation specifications related to the module that created these destinations.

Troubleshooting administration tasks and tools

Use the information in this group of topics to identify and resolve problems that can occur while you are administering the runtime environment.

Profile-specific log files

There are log files detailing the characteristics and runtime activities of individual profiles. These log files are located within the logs directory of the profile path.

There are a number of log files that are created for each profile. Some of these logs describe the parameters used for the creation of the profile. These types of log files generally remain unchanged once the profile is fully configured. Other profile-specific logs are continually updated to capture error, warning, and information messages emitted during runtime. Some of these log files are also used to capture a Common Base Event (that may include business object data) that is selected for monitoring.

The table below specifies the different types of profile-specific log files and the locations where you can find them within the product. Within the table, the variable *install_root* represents the installation directory of WebSphere Process Server. The variable *profile_root* represents the root location of a profile.

Table 1. Profile-specific log files updated during runtime

Log	Contents
First failure data capture (ffdc) log and exception files (common to all profile types) are found in <i>profile_root</i> /logs/ffdc	<p>Contains the ffdc log and exception files for individual profiles. There are two types of ffdc logs: a single log file with a compilation of all the errors encountered during the profile runtime, and numerous text files with details such as stack traces and other information. The naming conventions for the different types of profiles are given for both files, as follows:</p> <ul style="list-style-type: none">• Deployment manager profile:<ul style="list-style-type: none">– Log file — <i>deployment_manager_name_exception.log</i>.– Text files — <i>deployment_manager_name_hex_id_date_time.txt</i>.•<ul style="list-style-type: none">– Log file(s) — <i>node_agent_name_exception.log</i> and <i>server_name_exception.log</i>.– Text files — <i>node_agent_name(or)server_name_hex_id_date_time.txt</i>.• Managed node profile:<ul style="list-style-type: none">– Log file(s) — <i>node_agent_name_exception.log</i> and <i>server_name_exception.log</i>.– Text files — <i>node_agent_name(or)server_name_hex_id_date_time.txt</i>.• Stand-alone profile:<ul style="list-style-type: none">– Log file — <i>server_name_exception.log</i>.– Text files — <i>server_name_hex_id_date_time.txt</i>.

Table 1. Profile-specific log files updated during runtime (continued)

Log	Contents
<p>Deployment manager logs SystemErr.log and SystemOut.log are found in deployment manager job logs unless customized to write to a file. The start and stop server logs are found in <i>profile_root/logs/dmgr</i> if started in USS.</p>	<p>You will work primarily with two log files in this directory:</p> <ul style="list-style-type: none"> • startServer.log — Contains the system parameters detected on the system and the messages emitted by the deployment manager during the start process • stopServer.log — Contains the system parameters detected on the system and the messages emitted when the deployment manager is shut down.
<p>Node agent logs SystemErr.log and SystemOut.log are found in node agent job logs unless customized to write to a file. The start and stop server logs are found in <i>profile_root/logs/nodeagent</i> if started in USS.</p>	<p>You will work primarily with two log files in this directory:</p> <ul style="list-style-type: none"> • startServer.log — Contains the system parameters detected on the system and the messages emitted by the node agent during the start process • stopServer.log — Contains the system parameters detected on the system and the messages emitted when the node agent is shut down.
<p>Server logs SystemErr.log and SystemOut.log are found in server job logs unless customized to write to a file. The start and stop server logs are found in <i>profile_root/logs/<servername></i> if started in USS.</p>	<p>You will work primarily with two log files in this directory:</p> <ul style="list-style-type: none"> • startServer.log — Contains the system parameters detected on the system and the messages emitted by the server during the start process • stopServer.log — Contains the system parameters detected on the system and the messages emitted when the server is shut down.
<p>Node federation log files are found in the <i>profile_root/logs/</i> directory.</p>	<p>Two log files are generated when you attempt to federate a node to a deployment manager:</p> <ul style="list-style-type: none"> • addNode.log — contains the pertinent server environment information and messages generated when you attempt to federate the profile. • isFederated.log — lists the commands used by the deployment manager to federate the profile.
<p>The location of the Integrated Solutions Console application deployment log file is listed here (only for deployment manager and stand-alone profiles):The location of the Integrated Solutions Console application deployment log file is in the <i><profilepath>/logs</i> directory.</p>	<p>The <i>iscinstall.log</i> file contains information regarding the deployment of the administrative console application in a deployment manager or stand-alone profile.</p>
<p>The location of the log files for a profile creation is the <i>install_root/logs/</i> directory.</p>	<p>These files contain the output from profile creation. All profile types will contain this file.</p>

Related concepts

“WebSphere Process Server log files” on page 9

There are two distinct groups of log files in the installed product. Logs detailing the product installation, product updates, and profile management are one group. Logs detailing the characteristics and runtime activities of individual profiles make up the second group.

“Servers and recovery mode processing” on page 75

When you restart an application server instance with active transactions after a failure, the transaction service uses recovery logs to complete the recovery process.

Troubleshooting the failed event manager

This topic discusses problems that you can encounter while using the failed event manager.

Note: This topic does not discuss how to use the failed event manager to find, modify, resubmit, or delete failed events on the system. For information about managing failed events, see *Managing WebSphere Process Server failed events* in the information center.

Select the problem you are experiencing from the table below:

Problem	Refer to the following
I am having trouble entering values in the Search page's By Date tab	“Values in the By Date and From Date field automatically change to default if entered incorrectly”
I am having trouble deleting expired events	“Using the Delete Expired Events function appears to suspend the failed event manager” on page 22
I am having trouble with failed events not being created	“Failed events are not being created” on page 22

Values in the **By Date** and **From Date** field automatically change to default if entered incorrectly

The Search page's **From Date** and **To Date** fields require correctly formatted locale-dependent values. Any inconsistency in the value's format (for example, including four digits in the year instead of 2, or omitting the time) will cause the failed event manager to issue the following warning and substitute a default value in the field:

```
CWMAN0017E: The date entered could not be parsed correctly:  
your_incorrectly_formatted_date. Date: default_date is being used.
```

The default value of the **From Date** field is defined as January 1, 1970, 00:00:00 GMT.

Important: The actual default value shown in your failed event manager implementation will vary depending on your locale and time zone. For example, the From Date field defaults to 12/31/69 7:00 PM for a workstation with an en_US locale in the Eastern Standard Time (EST) time zone. The default value for the **To Date** field is always the current date and time, formatted for your locale and time zone.

To avoid this problem, always enter your dates and times carefully, following the example provided above each field.

Using the Delete Expired Events function appears to suspend the failed event manager

If you use the Delete Expired Events button in situations where there are many failed events in the current search results, or where those events contain a large amount of business data, the failed event manager can appear to be suspended indefinitely.

In this situation, the failed event manager is not suspended: it is working through the large data set, and will refresh the results set as soon as the command completes.

Failed events are not being created

If the Recovery subsystem is not creating failed events, go through the following checklist of potential causes:

- Ensure that the wpsFEMgr application is running. If necessary, restart it.
- Ensure that the failed event manager's database has been created, and that the connection has been tested.
- Ensure that the necessary failed event destination has been created on the SCA system bus. There should be one failed event destination for each deployment target.
- Ensure that the Quality of Service (QoS) **Reliability** qualifier has been set to Assured for any Service Component Architecture (SCA) implementation, interface, or partner reference that participates in events you want the Recovery service to handle.

Troubleshooting store-and-forward processing

This topic discusses problems that you can encounter with store-and-forward processing.

Select the problem you are experiencing from the table below:

Problem	Refer to the following
I am having problems setting the store-and-forward qualifier	"Store-and-forward qualifier processing only works on asynchronous interfaces" on page 23
Qualifying runtime exceptions are occurring, but events are not getting stored	"Store is not activated by qualifying runtime exceptions" on page 23
Messages are still being processed even though the Store and Forward widget shows the state is set to Store (Network deployment environment)	"In a network deployment environment, messages are being processed even though the store-and-forward state is set to Store" on page 23
The Store and Forward widget shows the state is set to Forward, but messages are not being processed by all members of the cluster. (Network deployment environment)	"In a network deployment environment, messages are not getting processed by all members of the cluster even though the store-and-forward state is set to Forward" on page 24

Store-and-forward qualifier processing only works on asynchronous interfaces

The store-and-forward qualifier must be specified on an asynchronous interface. The store cannot be activated if the interface is called synchronously.

Here are some guidelines (with respect to components) to help you determine if the interface is being called synchronously or asynchronously.

- Examine your short-running business process and what import it invokes. For example, JMS is an asynchronous import. Therefore, it is called asynchronously by a short-running process. HTTP is a synchronous import. Therefore, it is called synchronously.
- Long-running processes invoke imports based on the preferred interaction style set on the import's interface. Look at the interaction style set on the import's interface to see whether it is synchronous or asynchronous.

Note: You can find this setting on the interface's detail tab.

- POJO components invoke components based on the code that is written in the component. Look at the code written in the component to see whether it is synchronous or asynchronous.

Also, consider these restrictions:

- Store-and-forward qualifier cannot be set on long-running processes.
- Store-and-forward cannot be set on exports (except SCA export).

Store is not activated by qualifying runtime exceptions

If the store is not being activated by qualifying runtime exceptions, check the following.

- The exception specification in the store-and-forward qualifier matches the exception that occurs at runtime. If the exception specification does not match, storing does not activate.
- The user code in the path is not catching the exception and wrapping it. Or, it is converting it into a different exception. The exception received by the store-and-forward function can be viewed in the exception details for the failed event.
- The destination component for a failed event has a store-and-forward qualifier set on it. Storing is activated once a failed event is generated. If a failed event is generated for a component that is upstream from the component that has a store-and-forward qualifier set on it, then the store-and-forward component is being invoked synchronously and not asynchronously. If a failed event is generated for a component that is downstream from the store-and-forward qualifier component, rather than the component with the store-and-forward qualifier set on it, then there is an asynchronous invocation closer to the failure and the store-and-forward qualifier should be moved to that component.

In a network deployment environment, messages are being processed even though the store-and-forward state is set to Store

Messages might continue to be processed by some members of a cluster, despite the state being set to Store, if the state is not set to Store for each member of the

cluster. To fix this problem, confirm that the state is set to Store for each member of the cluster in the Store and Forward widget. If any members of the cluster are set to Forward, change them to Store.

This might also happen if one of the members of the cluster is forced to restart. Since the Store state is not persistent, it reverts to the Forward state at restart. To fix this problem, change the state to Store for the module in the Store and Forward widget.

Note: When the service becomes available again, you should not set the state to Store immediately if you want new events to be processed. If you set the state to Store before new events have the chance to be processed, they will be stored in the queue.

In a network deployment environment, messages are not getting processed by all members of the cluster even though the store-and-forward state is set to Forward

Messages might continue to be stored by some members of a cluster, despite the state being set to Forward, if the store-and-forward state is not set to Forward for each member of the cluster. To fix this problem, confirm that the state is set to Forward for the module in the Store and Forward widget. If any members of the cluster are set to Store, change them to Forward.

Note:

Troubleshooting the business rules manager

Some of the problems you might encounter using the business rules manager are login errors, login conflicts, and access conflicts.

You can take various steps to troubleshoot these problems.

Resolving login errors

A log in error occurs upon logging in.

Before you begin

About this task

The login error message is as follows:

Unable to process login. Please check User ID and password and try again.

Note: Login errors occur only when administrative security is enabled and either the user ID, password, or both, are incorrect.

To resolve login errors, perform the following steps.

Procedure

1. Click **OK** on the error message to return to the Login page.
2. Enter the valid **User ID** and **Password**.
 - If passwords are case sensitive, make sure that Caps Lock key is not on.
 - Make sure the user ID and password are spelled correctly.

- Check with the system administrator to be sure that the user ID and password are correct.
3. Click **Login**.

What to do next

If you resolve the login error, you will now be able to log in to the business rules manager. If the error is not resolved, contact your system administrator.

Resolving login conflict errors

A login conflict error occurs when another user with the same user ID is already logged in to the application.

Before you begin

About this task

The login conflict message is as follows:

Another user is currently logged in with the same User ID. Select from the following options:

Usually this error occurs when a user closed the browser without logging out. When this condition occurs, the next attempted login before the session timeout expires results in a login conflict.

Note: Login conflict errors occur only when administrative security is enabled.

To resolve login conflict errors, select from the following three options:

- Return to the Login page.
Use this option if you want to open the application with a different user ID.
- Log out the other user with the same user ID.
Use this option to log out the other user and start a new session.

Note: Any unpublished local changes made in the other session will be lost.

- Inherit the context of the other user with the same user ID and log out that user.
Use this option to continue work already in progress. All unpublished local changes in the previous session that have been saved will not be lost. The business rules manager will open to the last page displayed in the previous session.

Resolving access conflict errors

An access conflict error occurs when a business rule is updated in the data source by one user at the same time another user is updating the same rule.

Before you begin

This error is reported when you publish your local changes to the repository.

About this task

To correct access conflict errors, perform the following actions:

- Find the source of the business rule that is causing the error and check if your changes on the local machine are still valid. Your change may no longer be required after the changes done by another user.
- If you choose to continue working in the business rules manager, you must reload those business rule groups and rule schedules in error from the data source as your local changes of business rule groups and rule schedules in error are no longer usable. Reload a business rule group or rule schedule page, by clicking **Reload** in the Publish and Revert page of the rule for which the error was reported. You can still use local changes in other business rule groups and rule schedules that are not in error.

Troubleshooting WebSphere Application Server

Because IBM WebSphere Process Server is built on IBM WebSphere Application Server, the function that you are having problems with may be provided by the underlying WebSphere Application Server. You might want to consult troubleshooting information in the WebSphere Application Server documentation.

WebSphere Process Server is built on WebSphere Application Server, version 7.0.

For more information about troubleshooting in WebSphere Application Server, see the topic "Troubleshooting and support" in the WebSphere Application Server Information Center.

Related reference

 [Troubleshooting and support](#)

Tools for troubleshooting your applications

WebSphere Process Server and WebSphere Integration Developer include several tools you can use to troubleshoot your applications that you develop and deploy on the server.

During development of your applications, you can use debugging tools in WebSphere Integration Developer. You can implement runtime troubleshooting capabilities into your applications using logging, tracing, and service component event monitoring. Administrators of running applications can use the failed event manager to view, modify, resubmit, and delete failed operations between Service Component Architecture (SCA) components.

Debugging applications in WebSphere Integration Developer

In order to debug applications that are running on WebSphere Process Server, you must use your application development tool, such as IBM WebSphere Integration Developer.

About this task

For more information about debugging applications, see **Debugging components** in the IBM WebSphere Business Process Management information center or in the online documentation installed with WebSphere Integration Developer.

Related reference

 [IBM WebSphere Business Process Management information center](#)

Using logging, tracing, and monitoring in your applications

Designers and developers of applications that run on WebSphere Process Server can use capabilities such as monitoring and logging that add troubleshooting features to applications.

About this task

WebSphere Process Server is built on IBM WebSphere Application Server, Network Deployment. For more information, see the topic "Adding logging and tracing to your application" in the WebSphere Application Server Information Center.

To use logging, tracing, and monitoring with your applications, perform the following steps.

Procedure

- You can set up service component event monitoring for applications running on WebSphere Process Server. For more information, see the "Monitoring service component events" topic link in the Related Topics section at the bottom of this page.
- You can add logging and tracing to your applications using WebSphere Application Server.

Related reference

 [Adding logging and tracing to your application](#)

 [Monitoring service component events](#)

WebSphere Process Server monitoring can capture the data in a service component at a certain event point. You can view each event in a log file, or you can use the more versatile monitoring capabilities of a Common Event Infrastructure server.

Troubleshooting Service Component Architecture processing and call chains

Cross-Component Trace identifies whether a Service Component Architecture (SCA) operation completed successfully. It allows you to identify `systemout.log` or `trace.log` data that is associated with WebSphere Process Server and WebSphere Enterprise Service Bus modules and components. The log records associated with the WebSphere® ESB applications hold information about errors or events that occurred during processing and can be used for problem determination using WebSphere Integration Developer.

Events that can be captured include:

- Errors that occur during processing because of corrupted data.
- Errors when resources are not available, or are failing.
- Interpretation of code paths.

You can access the Cross-Component Trace page from the administrative console and then clicking **Troubleshooting** → **Cross-Component Trace**.

Handling and deleting collected data

Consider the following with regard to handling and deleting data collected by Cross-Component Trace:

- SCA call chain information is added to the `systemout.log` and `trace.log` files and is purged as those files are purged.
- Data snapshots capture the input and output data of call chains.

The input and output data is captured as files in the `logs\XCT` directory. To view this data using WebSphere Integration Developer, WebSphere Integration Developer needs access to the `systemout.log` files and the `logs\XCT` directory. If WebSphere Integration Developer is not available on the server, copying the logs directory and placing it on a machine (so that it can be accessed by WebSphere Integration Developer) preserves the file structure so that WebSphere Integration Developer can make use of the log files and the data snapshot files.

Note: WebSphere Integration Developer can use the data snapshot files where they are (without moving them) if it can access the files in the logs directory. If you need to move files, it is safest to move the entire logs directory. By moving the entire logs directory you get the XCT, first failure data capture (FFDC) files, and the `systemout.log` and `trace.log` files.

Data snapshot files are written to server-specific subdirectories using the following directory structure:

```
logs\  
  server  
  ffdc  
  xct\  
  
```

```
server-specific_dir\  
    2009-0-25-11  
    2009-0-26-12  
    2009-0-26-14
```

Where `server-specific_dir` name is derived from the name of the server. For example, `server1` is the default server name for a stand-alone installation.

- Data snapshot files in `logs\XCT\server` are referenced from the `systemout.log` and `trace.log` files that were created at the same time by the server. When WebSphere Application Server deletes the old `systemout.log` and `trace.log` files, the associated Cross-Component Trace data snapshot files in `logs\XCT\server` can also be deleted.

You can use the timestamps in the `systemout.log` and `trace.log` files to identify and determine what data snapshot files to delete. It is safe to delete all the data snapshot files for a server that are older than the oldest date in the `systemout.log` and `trace.log` files. Preferably, you should use the **Delete data snapshot files** function from the administrative console when data snapshot files are no longer needed. For detailed information on the ways that you can delete data snapshot files, see *Deleting data snapshot files*.

- Do not save or add files to the `logs\XCT` directory. Do not copy or create new directories into the `logs\XCT` directory.

WebSphere Process Server manages the content of the `logs\XCT` directory and deletes items that are no longer needed. WebSphere Process Server might consider unrecognized files or directories as unnecessary and delete them. If you want to save a copy of the data snapshot files, copy the data to another directory outside of the `logs\XCT` directory.

Cross-Component Trace settings and call chain processing

The information in this section describes the affect that Cross-Component Trace configuration settings have on call-chain processing.

The information in this section includes a description of various Cross-Component Trace configurations and explains the call chain events that result from the configurations.

General rules on call chain processing and Cross-Component Trace configuration decisions

- If Cross-Component Trace is turned off for a server, then no Cross-Component Trace records are written to that server's logs.
- Cross-Component Trace configuration settings for a particular server, *affect that server only*.

For example, if Server A has **Trace all = Yes** and Server B has **Trace all = No**, Cross-Component call chains are in the logs for Server A only. Similarly, this rule applies to setting the data snapshot feature. If **Enable data snapshot = Yes** on Server A and **Enable data snapshot = No** on Server B, then only Server A will have data snapshot files in its logs directory.

- Application-specific Cross-Component Trace data flows between servers that have the **Enable Cross-Component Trace = Yes**.

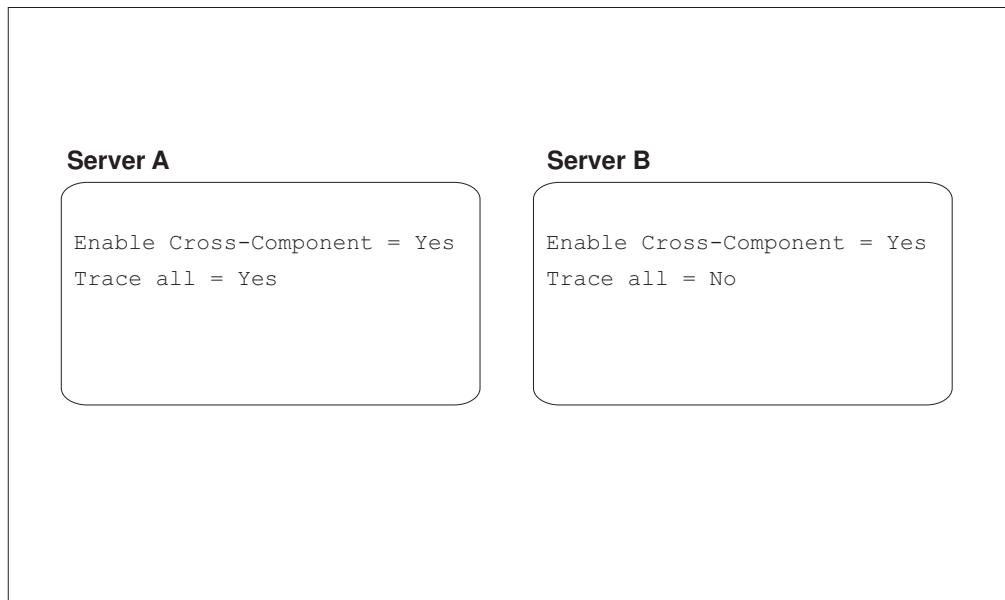
For example, if both Server A and Server B have **Enable Cross-Component Trace = Yes** and Server A has enabled Cross-Component Trace for a specific SCA module, the calls made from the Cross-Component Trace-enabled module on Server A (to applications or services on Server B), will result in Server A having call chains for all of activity related to the Cross-Component Trace-enabled

module. Server B would also have call chains, but only for those calls that came from the Cross-Component Trace-enabled module on Server A. The logs of the two servers can be combined to reveal the entire call chain activity.

- To create a Cross-Component Trace for long running BPEL process instances, you must select **Enable Cross-Component Trace** and **Trace all**, or enable Cross-Component Trace for the desired SCA module before the BPEL process instance is created.

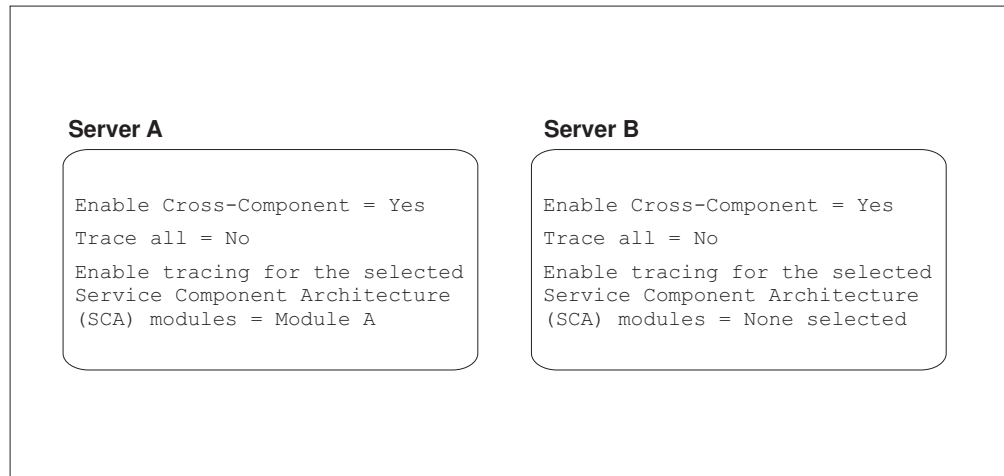
For more information see *Enabling Cross-Component Trace for BPEL long-running instances* in the Cross-Component Trace overview.

The following illustration is of two servers (Server A and Server B), both with Cross-Component Trace enabled. Server A has the **Trace all** value set to Yes, while Server B has **Trace all** set to No.



Result: For the Cross-Component Trace configuration scenario illustrated above, call chain events would result on Server A, but not on Server B.

The following illustration is of two servers (Server A and Server B), both with Cross-Component Trace enabled. Server A has the **Trace all** value set to No and it includes Module A as a module on which to enable Cross-Component Trace. Server B has **Trace all** set to No and has no SCA modules selected for Cross-Component Trace.



Result: For the Cross-Component Trace configuration scenario illustrated above, call chain events would result on Server A. Trace activity for all Module A operations are written to the log on Server A. Any calls made from Module A to applications or services on Server B, results in call chains. The call chains on Server B would only pertain to those calls that came from Module A (because that module is configured for Cross-Component Trace).

Managing failed events

The WebSphere Process Server Recovery service captures data about failed events. You can then use the failed event manager to view, modify, resubmit, or delete the failed event.

The WebSphere Process Server Recovery service manages failed operations between Service Component Architecture (SCA) components, failed JMS events, failed WebSphere MQ events, and failed operations within long-running business processes.

Note: For service runtime exceptions that are generated when a requested service is unavailable, you can use the Store and Forward feature to prevent further failures. You specify a store-and-forward qualifier when you configure a component that will be invoked asynchronously. When a runtime error is generated by that component, subsequent events (in this case, asynchronous requests) are prevented from reaching the component. See “Preventing failures when a service is unavailable” for more information.

Failed SCA events

In the context of SCA, an event is a request or response that is received by a service application. It can come from an external source (such as an inbound application adapter) or an external invocation to a Web service. The event consists of a reference to the business logic it wants to operate and its data, stored in a Service Data Object (a business object). When an event is received, it is processed by the appropriate application business logic.

A single thread of execution can branch off into multiple branches (or threads); the individual branches are linked to the main invoking event by the same session context.

If this business logic in one of these branches cannot execute completely due to system failure, component failure, or component unavailability, the event moves into the failed state. If multiple branches fail, a failed event is created for each. The Recovery service handles the following types of failed SCA events:

- Event failures that occur during an asynchronous invocation of an SCA operation
- Event failures that are caused by a runtime exception (for example, any exception that is not declared in the methods used by the business logic)

The Recovery service does not handle failures from synchronous invocations.

Failed SCA events typically have source and destination information associated with them. The source and destination are based on the failure point (the location where the invocation fails), regardless of the type of interaction. Consider the following example, where Component A is asynchronously invoking Component B. The request message is sent from A to B, and the response (callback) message is sent from B to A.

- If the exception occurs during the initial request, Component A is the source and Component B is the destination for the purposes of the failed event manager.
- If the exception occurs during the response, Component B is the source and Component A is the destination for the purposes of the failed event manager.

This is true for all asynchronous invocations.

The Recovery service sends failed SCA asynchronous interactions to failed event destinations that have been created on the SCA system bus (SCA.SYSTEM.cell_name.Bus). The data for failed events is stored in the failed event database (by default, WPCRSDB) and is made available for administrative purposes through the failed event manager interface.

Failed WebSphere MQ events

A WebSphere MQ event can fail when there is a problem (such as a data-handling exception) in the WebSphere MQ binding export or import used by an SCA module.

WebSphere Integration Developer provides a recovery binding property that allows you to enable or disable recovery for each WebSphere MQ binding at authoring time. The recoveryMode property can be set to one of the following:

bindingManaged	Allow binding to manage recovery for failed messages
unmanaged	Rely on transport-specific recovery for failed messages

Recovery for WebSphere MQ bindings is enabled by default. When it is enabled, WebSphere MQ failed events are created in the following situations:

- The function selector fails
- The fault selector fails
- The fault selector returns the RuntimeException fault type
- The fault handler fails
- The data binding or data handler fails after a single retry in WebSphere MQ

In addition, a failed SCA event is created when the `ServiceRuntimeException` exception is thrown in a WebSphere MQ binding target component after a single retry in WebSphere MQ.

These failures can occur during inbound or outbound communication. During outbound communication, MQImport sends a request message and receives the response message; a failed event is generated if the WebSphere MQ import binding detects a problem while processing the service response. During inbound communication, the sequence of events is as follows:

1. MQExport receives the request message.
2. MQExport invokes the SCA component.
3. The SCA component returns a response to MQExport.
4. MQExport sends a response message.

A failed event is generated if the WebSphere MQ export binding detects a problem while processing the service request.

The Recovery service captures the WebSphere MQ message and stores it in the failed event database. It also captures and stores the module name, component name, operation name, failure time, exception detail, and WebSphere MQ properties of the failed event. You can use the failed event manager or a custom program to manage failed WebSphere MQ events, including resubmitting or deleting the event.

To disable recovery, you must explicitly disable it in WebSphere Integration Developer by setting the `recoveryMode` property to `unmanaged`.

Note: If the `recoveryMode` property is missing (for earlier versions of applications), the recovery capability is regarded as enabled. When recovery is disabled, a failed message is rolled back to its original destination and retried. The system does not create a failed event.

Failed JMS events

The Java Message Service (JMS) binding type and configuration determine whether a failed event is generated and sent to the failed event manager.

WebSphere Integration Developer provides a recovery binding property that allows you to enable or disable recovery for each JMS binding at authoring time. The `recoveryMode` property can be set to one of the following:

<code>bindingManaged</code>	Allow binding to manage recovery for failed messages
<code>unmanaged</code>	Rely on transport-specific recovery for failed messages

Recovery for JMS bindings is enabled by default. When it is enabled, JMS failed events are created in the following situations:

- The function selector fails
- The fault selector fails
- The fault selector returns the `RuntimeException` fault type
- The fault handler fails
- The data binding or data handler fails after a single retry in JMS

In addition, a failed SCA event is created when the `ServiceRuntimeException` exception is thrown in a JMS binding target component after a single retry in JMS.

These failures can occur during inbound or outbound communication. During outbound communication, `JMSImport` sends a request message and receives the response message; a failed event is generated if the JMS import binding detects a problem while processing the service response. During inbound communication, the sequence of events is as follows:

1. `JMSEExport` receives the request message.
2. `JMSEExport` invokes the SCA component.
3. The SCA component returns a response to `JMSEExport`.
4. `JMSEExport` sends a response message.

A failed event is generated if the JMS export binding detects a problem while processing the service request.

The Recovery service captures the JMS message and stores it in a Recovery table in the Common database. It also captures and stores the module name, component name, operation name, failure time, exception detail, and JMS properties of the failed event. You can use the failed event manager to manage failed JMS events, or you can use a custom program.

To disable recovery, you must explicitly disable it in WebSphere Integration Developer by setting the `recoveryMode` property to `unmanaged`.

Note: If the `recoveryMode` property is missing (for earlier versions of applications), the recovery capability is regarded as enabled. When recovery is disabled, a failed message is rolled back to its original destination and retried. The system does not create a failed event.

Failed Business Process Choreographer events

In the context of Business Process Choreographer, exceptions can occur that, if not handled by the process logic, cause an activity to stop or the process instance to fail. A failed event is generated when a long-running Business Process Execution Language (BPEL) process fails and one of the following happens:

- The process instance enters the failed or terminated state
- An activity enters the stopped state

The Recovery service captures the module name and component name for failed Business Process Choreographer events. Failed event data is stored in the Business Process Choreographer database (BPEDB) database.

Note that the Recovery service does not handle failures from business process and human task asynchronous request/reply invocations.

Business Flow Manager hold queue messages

You can use the failed event manager to manage navigation messages that are stored in the Business Flow Manager hold queue. A navigation message might be stored in the hold queue if:

- An infrastructure, such as a database, is unavailable.
- The message is damaged.

In a long-running process, the Business Flow Manager can send itself request messages that trigger follow-on navigation. These messages trigger either a process-related action (for example, invoking a fault handler) or an activity-related action (for example, continuing process navigation at the activity). A navigation message always contains its associated process instance ID (piid). If the message triggers an activity-related action, it also contains the activity template ID (atid) and the activity instance ID (aiid).

You can use the failed event manager to manage Business Flow Manager hold queue messages, or you can use a custom program.

Business Flow Manager hold queue messages cannot be deleted directly in the failed event manager. If the related process instance does not exist, replaying the hold queue message will result in deletion of the message.

How are failed events managed?

An administrator uses the failed event manager to browse and manage failed events. Common tasks for managing failed events include:

- Browsing all failed events
- Searching for failed events by specific criteria
- Editing data for a failed event
- Resubmitting failed events
- Deleting failed events

To access the failed event manager, click **Integration Applications** → **Failed Event Manager**.

Related concepts

“Retention queues and hold queues” on page 76

When a problem occurs while processing a message, it is moved to the retention queue or hold queue.

Security considerations for recovery

If you have enabled security for your WebSphere Process Server applications and environment, it is important to understand how role-based access and user identity affect the Recovery subsystem.

Role-based access for the failed event manager

The failed event manager uses role-based access control for the failed event data and tasks. Only the administrator and operator roles are authorized to perform tasks within the failed event manager. Users logged in as either administrator or operator can view all data associated with failed events and can perform all tasks.

Event identity and user permissions

A failed event encapsulates information about the user who originated the request. If a failed event is resubmitted, its identity information is updated to reflect the user who resubmitted the event. Because different users logged in as administrator or operator can resubmit events, these users must be given permissions to the downstream components required to process the event.

For more information about implementing security, see *Securing applications and their environment*.

Finding failed events

Failed events are stored in a database and are retrieved through the search functionality of the failed event manager. You can search for all failed events on all the servers within the cell, or for a specific subset of events.

Before you begin

If administrative security is enabled, you must be logged in as administrator or operator to perform this task.

About this task

This topic describes how to find all failed events in the cell. This default query returns all SCA and JMS failed events.

If Business Process Choreographer is installed, the query also returns failed, terminated, and stopped Business Process Choreographer events.

To retrieve a complete list of failed events, use the following procedure.

Procedure

1. Ensure the administrative console is running.
2. Click **Integration Applications** → **Failed Event Manager** to enter the failed event manager.
3. From the **Failed events on this server** box, click **Get all failed events**.

Results

The Search Results page opens, displaying a list of all the WebSphere Process Server failed events in the cell.

What to do next

You can now view (and in some cases, modify) data in a failed event, resubmit it, or delete it.

Searching for events by criteria

Use the failed event manager Search page to find only those events that meet specified criteria. You can search by failed event type and by criteria such as failure time, event destination or source, exception or business object type, session ID or event sequencing qualifier.

Before you begin

If administrative security is enabled, you must be logged in as administrator or operator to perform this task.

About this task

To search for a specific subset of failed events on the server, perform the following steps.

Procedure

1. Ensure the administrative console is running.

2. Click **Integration Applications** → **Failed Event Manager** to enter the failed event manager.
3. From the **Failed events on this server** box, click **Search failed events**.
4. From the **Event type** box on the Search failed events page, select one or more types of events to search for:
 - SCA
 - JMS
 - WebSphere MQ
 - Business Process Choreographer
 - Business Flow Manager hold queue messages
5. If you are searching for Business Process Choreographer events, verify the event status selected in the Event status box. By default, the failed event manager returns all failed, stopped, and terminated Business Process Choreographer events, but you can modify the search to return only events with a particular status.
6. Optional: Specify any additional search criteria. The following table describes the available options. If you specify multiple criteria, the AND operator is used during the query; the failed event manager returns only events that meet all of the criteria.

Table 2. Search criteria

Search criteria	Field or fields to use	Supported event types	Usage notes
The module, component, or method the event was en route to when it failed.	Module Component Operation	SCA JMS WebSphere MQ Business Process Choreographer Business Flow Manager hold queue	Use one or more of these fields to search for failed events associated with a specific module, component, or method.
The time period during which the event failed	From date To date	SCA JMS WebSphere MQ Business Process Choreographer Business Flow Manager hold queue	Formats for date and time are locale-specific. An example is provided with each field. If the value you provide is not formatted correctly, the failed event manager displays a warning and substitutes the default value for that field. The time is always local to the server. It is not updated to reflect the local time of the individual workstations running the administrative console.

Table 2. Search criteria (continued)

Search criteria	Field or fields to use	Supported event types	Usage notes
The session in which the event failed	Session ID	SCA	None
The module or component from which the event originated	Source module Source component	SCA	Use one or both of these fields to find only those failed events that originated from a specific source module or component. The failed event manager determines the source based on the point of failure, regardless of interaction type.
The type of business object in the failed event	Business object type	SCA	None
Whether the event had the event sequencing qualifier specified	Event sequencing qualifier	SCA	None
Whether the event caused the store to be started	Store and forward qualifier	SCA Business Process Choreographer	None
Whether the event was caused because a failure response could not be sent to Business Process Choreographer	Process response qualifier	SCA	None
The exception thrown when the event failed	Exception text	SCA	Specify all or part of the exception text in the field to find all events associated with that exception.

For detailed information about each field and the values it accepts, see the online help for the failed event manager Search page.

- Click **OK** to begin the search.

What to do next

You can now view (and in some cases, modify) data in a failed event, resubmit it, or delete it.

Working with data in failed events

Each failed event has data associated with it; often, that data can be edited before an event is resubmitted. There are two basic types of data for a failed event: data about the event, and business data.

Data about the failed event

All failed events have the following data:

- The event ID, type, and status
- The time the event failed
- The deployment target associated with the event

In addition, SCA, JMS, WebSphere MQ, Business Process Choreographer, and Business Flow Manager hold queue events have data specific to the event type:

- SCA events
 - The session ID
 - The type of service invocation used between SCA components
 - The names of the module and component from which the event originated (the source).
 - The names of the destination module, component and method for the event
 - Whether an event sequencing qualifier has been declared for this event
 - The destination module where the event has been or will be resubmitted
 - The correlation ID, if one exists
 - The exception thrown when the event failed
 - The expiration date for resubmitted events (this data can be edited)
 - The trace control set for the event (this data can be edited)
- JMS events:
 - The type of service invocation used
 - The names of the destination module, component and method for the event
 - The exception thrown when the event failed
 - The destination module where the event has been or will be resubmitted
 - The correlation ID, if one exists
 - The expiration date for resubmitted events (this data can be edited)
 - The JMS-specific properties associated with the failed event:
 - The message type and priority
 - The JMS destination
 - The delivery mode
 - Redelivery data, including the redelivered count and redelivered indicator (true or false)
 - The destination replies are sent to for request-response or two-way interactions
- WebSphere MQ events:
 - The type of service invocation used
 - The names of the destination module, component and method for the event
 - The exception thrown when the event failed
 - The destination module where the event has been or will be resubmitted
 - The correlation ID, if one exists
 - The expiration date for resubmitted events (this data can be edited)
 - The WebSphere MQ-specific properties associated with the failed event:
 - The message type, format, and priority
 - The WebSphere MQ destination
 - The delivery mode

- Redelivery data, including the redelivered count and redelivered indicator (true or false)
- The reply-to queue and queue manager
- Business Process Choreographer events:
 - The names of the destination module and component for the event
 - The process instance name associated with the event
 - The top-level process ID associated with the event
- Business Flow Manager hold queue events:
 - The process instance ID (if the process instance does not exist, 0 is returned)
 - The name and state of the process instance
 - The name of the associated process template
 - The activity instance name and ID
 - The activity template ID

Business data

SCA and Business Process Choreographer failed events typically include business data. Business data can be encapsulated in a business object, or it can be simple data that is not part of a business object. Business data for SCA failed events can be edited with the business data editor available in the failed event manager.

Browsing data in failed events

Use the failed event manager to view failed event data and any business data associated with the event.

Before you begin

If administrative security is enabled, you must be logged as administrator or operator to perform this task.

About this task

To browse failed event data, use the following procedure.

Procedure

1. Ensure that the failed event manager is open and that you have retrieved a list of the failed events on your system.
2. From the Search Results page of the failed event manager, click the ID (found in the Event ID column) of the failed event whose data you want to browse. The Failed Event Details page opens and displays all of the information about the event.
3. If your failed event has business data, you can browse it by clicking **Edit business data**.

The Business Data Editor collection page opens, displaying the business data associated with the failed event. Each parameter name in the hierarchy is a link. If the parameter is a simple data type, clicking its name will open up a form so you can edit the parameter's value. If the parameter is a complex data type, clicking its name will expand the hierarchy further.

Editing trace or expiration data in a failed SCA event

The Failed Event Details page enables you to set or modify values for the trace control and expiration date associated with a failed event.

Before you begin

You must be logged in as administrator or operator to perform this task.

About this task

Important: Any edits you make to the trace or expiration data are only saved locally until you resubmit the event. If you perform any other action before resubmitting the event, all edits are lost.

Failed Service Component Architecture (SCA) events can be resubmitted with trace to help you monitor the event processing. Tracing can be set for a service or a component, and it can be sent to a log or to the Common Event Infrastructure (CEI) server. When you view the failed event data on the Failed Event Details page, the default trace value `SCA.LOG.INFO;COMP.LOG.INFO` is shown for the event. If you resubmit the event with this default setting, no trace occurs when the session calls an SCA service or executes a component.

Some failed SCA events also have an expiration. If a user has specified an expiration with the asynchronous call that sends the event, that data persists even if the event fails, and the expiration time appears in the **Resubmit Expiration Time** field on the Failed Event Details page. Expired failed events cannot be resubmitted successfully. To prevent a second failure, you can edit the expiration date for the event to ensure that it is not expired when it is resubmitted.

To edit trace or expiration data in a failed event, use the following procedure.

Procedure

1. Ensure that the failed event manager is open and that you have retrieved a list of the failed events on your system.
2. From the failed event manager's Search Results page, click the ID (found in the Event ID column) of the failed event whose data you want to edit.
The Failed Event Details page opens.
3. If the event has an expiration date that causes it to expire before it is resubmitted, edit the expiration in the **Resubmit expiration time** field.
The expiration time shown is local to the server. The value for this field must be formatted according to your specified locale. An example of the correct format for your locale is provided above the field.
4. If you want to enable tracing for the failed event, specify a new value in the **Trace Control** field. For detailed information about trace values, see the Monitoring topics in the WebSphere Business Process Management Information Center.
5. Do one of the following:
 - If the edited data is correct and you want to resubmit the event, click **Resubmit** to make the changes at a server level.
 - If you want to remove the changes you made, click **Undo local changes**.The edited failed event is resubmitted for processing and is removed from the failed event manager.

Related tasks

“Finding failed events” on page 38

Failed events are stored in a database and are retrieved through the search functionality of the failed event manager. You can search for all failed events on all the servers within the cell, or for a specific subset of events.

Editing business data in a failed SCA event

Business data can be encapsulated into a business object, or it can be simple data that is not part of a business object. A failed event can have both simple data and a business object associated with it. Use the business data editor to edit the business data associated with a failed event before you resubmit it.

Before you begin

If administrative security is enabled, you must be logged in as administrator or operator to perform this task.

About this task

For each failed event, the editor displays the associated business data in a hierarchical format; the navigation tree at the top of the table is updated as you navigate through the parameters to give you a clear picture of where you are in the hierarchy.

You can edit only simple data types (for example, String, Long, Integer, Date, Boolean). If a data type is complex (for example, an array or a business object), you must navigate through the business data hierarchy until you reach the simple data types that make up the array or business object. Complex data is denoted by an ellipsis (...) in the Parameter Value column.

Note you cannot use the failed event manager to edit business data for a Business Process Choreographer event. Instead, click the **Open calling process in Business Process Choreographer Explorer** link from the failed event detail page and use Business Process Choreographer Explorer to make any permitted modifications.

Important: Any edits you make to business data are saved locally. Changes are not made to the corresponding business data in the server until you resubmit the failed event.

To edit business data associated with a failed Service Component Architecture (SCA) event, use the following procedure.

Procedure

1. Ensure that the failed event manager is open and that you have retrieved a list of the failed events on your system.
2. From the failed event manager's Search Results page, click the ID (found in the Event ID column) of the failed event that has data you want to edit.
3. From the failed event detail page, click **Edit business data** to access the Business Data Editor collection page.
This page displays a hierarchical view of all the data associated with the failed event.
4. Navigate through the business data hierarchy by clicking the name of each parameter (these appear as links in the Parameter Name column). When you have located the parameter with values you want to edit, click its name.

If the parameter has an editable value, the Business Data Editor page opens.

5. In the **Parameter value** field, specify the new value for the parameter.

6. Click **OK**.

The change is saved locally and you are returned to the Business Data Editor collection page.

7. If you want to remove the changes you made, click **Undo local business data changes**.

All the edits are removed and the business data is returned to its original state.

8. If the edited business data is correct, click **Resubmit** to make changes at a server level.

The edited failed event is resubmitted for processing and is removed from the failed event manager.

Resubmitting failed events

If you want to send an event again, you must resubmit it from the failed event manager. You can resubmit an event without changes, or, in some cases, you can edit the business data parameters before resubmitting it.

When a failed event is resubmitted, the processing resumes only for the failed branch, not for the entire event.

Tracing is available for resubmitted SCA events to help monitor the event's processing. Tracing can be set for a service or a component, and its output can be sent to a log or to the Common Event Infrastructure (CEI) server.

You can also use the event's unique event ID to track its success or failure. If a resubmitted event fails again, it is returned to the failed event manager with its original event ID and an updated failure time.

Resubmitting an unchanged failed event

You can resubmit one or more unchanged failed events to be processed again. Processing resumes only for the failed branch, not for the entire event.

About this task

If administrative security is enabled, you must be logged in as administrator or operator to perform this task.

Procedure

1. Ensure that the failed event manager is open and that you have retrieved a list of the failed events on your system.
2. From the Search Results page, select the check box next to each failed event you want to resubmit.
3. Click **Resubmit**.

Results

Each selected event is resubmitted for processing and is removed from the failed event manager.

Resubmitting a failed SCA event with trace

You can monitor the resubmission of a failed Service Component Architecture (SCA) event to determine whether it now succeeds. The failed event manager provides optional tracing for all failed events.

About this task

Tracing can be set for a service or a component, and it can be output to a log or to the Common Event Infrastructure (CEI) server. For detailed information about setting and viewing trace, see the Monitoring topics in the information center.

If administrative security is enabled, you must be logged in as administrator or operator to perform this task.

Procedure

1. Ensure that the failed event manager is open and that you have retrieved a list of the failed events on your system.
2. From the Search Results page, select the check box next to each failed event you want to resubmit.
3. Click **Resubmit with trace**.
4. From the Resubmit with Trace page, specify the level of trace you want to use in the **Trace control** field.
By default, the value is `SCA.LOG.INFO;COMP.LOG.INFO`. With this setting, no trace occurs when the session calls an SCA service or executes a component.
5. Click **OK** to resubmit the failed event and return to the Search Results page.

What to do next

To view the trace log for a resubmitted event, open the corresponding component logger or use the CEI log viewer.

Resubmitting failed Business Process Choreographer responses

When a failure response cannot be delivered to the requesting business process due to an infrastructure problem, an event is stored in the failed event database. These type of events will have a process response qualifier specified on them. You can resubmit these failed events to either the request queue or the response queue using the failed event manager.

About this task

To resubmit a failed SCA event, perform the following steps.

Procedure

1. Ensure that the failed event manager is open and that you have retrieved a list of the failed events on your system.
2. From the Search Results page, select the check box next to each failed event you want to resubmit.
3. Click **Resubmit** or **Resubmit with trace**.
4. If the Process Response event qualifier is defined for the failed event, a resubmit page will appear. Select **Resubmit requests to the destination** or **Resubmit the exception response to the source**. Selecting **Resubmit the exception response to the source** allows the event to be sent to the response queue without having to be reprocessed.

Results

Depending on whether you selected **Resubmit requests to the destination** or **Resubmit the exception response to the source**, the event will be resubmitted to the appropriate queue.

Managing failed SCA events

When problems processing a Service Component Architecture (SCA) request or response message create a failed SCA event in the Recovery subsystem, you must decide how to manage that event. Use the information in this topic to help you identify and fix the error and clear the event from the Recovery subsystem.

About this task

Failed SCA events typically have source and destination information associated with them. The source and destination are based on the failure point (the location where the invocation fails), regardless of the type of interaction. Because runtime exceptions are not declared as part of the interface, component developers should attempt to resolve the exception and thus prevent a runtime exception from inadvertently being propagated to the client if the client is a user interface.

To manage a failed SCA event, perform the following steps.

Procedure

1. Use the failed event manager to locate information about the failed SCA event, taking note of the exception type.
2. Locate the exception type in Table 3 to determine the location and possible causes of the error, as well as suggested actions for managing the failed event.

Table 3. Failed SCA events

Exception type	Possible cause of error	Suggested action
ServiceBusinessException	A business exception occurred during the execution of a business operation.	Look at the exception text to determine the exact cause, and then take appropriate action.
ServiceExpirationRuntimeException	A SCA asynchronous message has expired.	Set the expiration time using the RequestExpiration qualifier on the service reference. Investigate why the service is not responding fast enough.
ServiceRuntimeException	A runtime exception occurred during the invocation or execution of a service.	Look at the exception text to determine the exact cause, and then take appropriate action.
ServiceTimeoutRuntimeException	Response to an asynchronous request was not received within the configured period of time.	Set the expiration time using the RequestExpiration qualifier on the service reference. Investigate why the service is not responding fast enough.
ServiceUnavailableException	This exception is used to indicate that there was an exception thrown while invoking an external service via an import.	Look at the exception text to determine the exact cause, and then take appropriate action.
ServiceUnwiredReferenceRuntimeException	A SCA reference used to invoke a service is not wired correctly.	Look at the exception text to determine the exact cause, and then take appropriate action to correctly wire the SCA reference.

Managing failed JMS events

When problems processing a JMS request or response message create a failed JMS event in the Recovery subsystem, you must decide how to manage that event. Use the information in this topic to help you identify and fix the error and clear the event from the Recovery subsystem.

About this task

To manage a failed JMS event, perform the following steps.

Procedure

1. Use the failed event manager to locate information about the failed JMS event, taking note of the exception type.
2. Locate the exception type in Table 4 to determine the location and possible causes of the error, as well as suggested actions for managing the failed event.

Table 4. Failed JMS events

Exception type	Location of error	Possible cause of error	Suggested action
FaultServiceException	Fault handler or fault selector	There is malformed data in the JMS message.	<ol style="list-style-type: none">1. Inspect the JMS message and locate the malformed data.2. Repair the client that originated the message so it creates correctly formed data.3. Resend the message.4. Delete the failed event.
		There was an unexpected error in the fault handler or fault selector.	<ol style="list-style-type: none">1. Debug the custom fault selector or fault handler, fixing any errors identified.2. Resubmit the failed event.
ServiceRuntimeException	Fault handler	The fault selector and runtime exception handler are configured to interpret the JMS message as a runtime exception. This is an expected exception.	Look at the exception text to determine the exact cause, and then take appropriate action.

Table 4. Failed JMS events (continued)

Exception type	Location of error	Possible cause of error	Suggested action
DataBindingException or DataHandlerException	Data binding or data handler	There is malformed data in the JMS message.	<ol style="list-style-type: none"> 1. Inspect the JMS message and locate the malformed data. 2. Repair the client that originated the message so it creates correctly formed data. 3. Resend the message. 4. Delete the failed event.
		There was an unexpected error in the data binding or data handler.	<ol style="list-style-type: none"> 1. Debug the custom data binding or data handler, fixing any errors identified. 2. Resend the message. 3. Delete the failed event.
SelectorException	Function selector	There is malformed data in the JMS message.	<ol style="list-style-type: none"> 1. Inspect the JMS message and locate the malformed data. 2. Repair the client that originated the message so it creates correctly formed data. 3. Resend the message. 4. Delete the failed event.
		There was an unexpected error in the function selector.	<ol style="list-style-type: none"> 1. Debug the custom function selector, fixing any errors identified. 2. Resend the message. 3. Delete the failed event.

Managing failed WebSphere MQ events

When problems processing a WebSphere MQ request or response message create a failed WebSphere MQ event in the Recovery subsystem, you must decide how to manage that event. Use the information in this topic to help you identify and fix the error and clear the event from the Recovery subsystem.

About this task

To manage a failed WebSphere MQ event, perform the following steps.

Procedure

1. Use the failed event manager to locate information about the failed event, taking note of the exception type.
2. Locate the exception type in Table 5 on page 50 to determine the location and possible causes of the error, as well as suggested actions for managing the failed event.

Table 5. Failed WebSphere MQ events

Exception type	Location of error	Possible cause of error	Suggested action
FaultServiceException	Fault handler or fault selector	There is malformed data in the WebSphere MQ message.	<ol style="list-style-type: none"> 1. Inspect the message and locate the malformed data. 2. Repair the client that originated the message so it creates correctly formed data. 3. Resend the message. 4. Delete the failed event.
		There was an unexpected error in the fault handler or fault selector.	<ol style="list-style-type: none"> 1. Debug the custom fault selector or fault handler, fixing any errors identified. 2. Resubmit the failed event.
ServiceRuntimeException	Fault handler	The fault selector and runtime exception handler are configured to interpret the WebSphere MQ message as a runtime exception. This is an expected exception.	Look at the exception text to determine the exact cause, and then take appropriate action.
DataBindingException or DataHandlerException	Data binding or data handler	There is malformed data in the WebSphere MQ message.	<ol style="list-style-type: none"> 1. Inspect the message and locate the malformed data. 2. Repair the client that originated the message so it creates correctly formed data. 3. Resend the message. 4. Delete the failed event.
		There was an unexpected error in the data binding or data handler.	<ol style="list-style-type: none"> 1. Debug the custom data binding or data handler, fixing any errors identified. 2. Resend the message. 3. Delete the failed event.

Table 5. Failed WebSphere MQ events (continued)

Exception type	Location of error	Possible cause of error	Suggested action
SelectorException	Function selector	There is malformed data in the WebSphere MQ message.	<ol style="list-style-type: none"> 1. Inspect the message and locate the malformed data. 2. Repair the client that originated the message so it creates correctly formed data. 3. Resend the message. 4. Delete the failed event.
		There was an unexpected error in the function selector.	<ol style="list-style-type: none"> 1. Debug the custom function selector, fixing any errors identified. 2. Resend the message. 3. Delete the failed event.

Managing stopped Business Process Choreographer events

Use the failed event manager and Business Process Choreographer Explorer to manage stopped Business Process Choreographer events in any process state. Stopped events occur if a Business Process Execution Language (BPEL) instance encounters an exception and one or more activities enter the Stopped state.

About this task

You can view, compensate, or terminate the process instance associated with a stopped Business Process Choreographer event. In addition, you can work with the activities associated with the event. viewing, modifying, retrying, or completing them as appropriate.

To manage stopped events originating from a long-running BPEL process, perform the following steps.

Procedure

1. Ensure the administrative console is running.
2. Open the failed event manager by clicking **Integration Applications** → **Failed Event Manager**.
3. Perform a search to find the stopped Business Process Choreographer event or events you want to manage.
4. For each stopped event you want to manage, do the following:
 - a. Click the stopped event ID in the Event ID column of the Search Results page.
 - b. From the event detail page, click **Open calling process in Business Process Choreographer Explorer**.
 - c. Use Business Process Choreographer Explorer to manage the event and its associated activities.

Finding business process instances related to a failed event

If a failed event is generated from a business process, the failed event manager provides a link to view that business process instance in Business Process Choreographer Explorer.

Before you begin

You must be logged in as administrator or operator to perform this task.

About this task

Examining the business process instance that generated the failed event can give you additional information about how or why the event failed. The business process instance and the failed event are linked by a common session ID.

Note: Not all failed events are generated from a business process instance.

To find and examine a business process instance related to a failed event, use the following procedure.

Procedure

1. From within the administrative console, use the failed event manager to locate the failed event you want to investigate. See “Finding failed events” on page 38 for instructions on how to search for failed events.
2. From the Failed Event Details page for that event, click **Open calling process in Business Process Choreographer Explorer**.

Results

The Business Process Choreographer Explorer opens in a new browser window and displays information about the related process instance.

Finding Common Base Events related to a failed event

A failed event can be related to one or more Common Base Events. The failed event manager provides a link to view related Common Base Events in the Common Base Event Browser.

Before you begin

You must be logged in as administrator or operator to perform this task.

About this task

Examining related Common Base Events can give you additional information about how or why the original event failed. The failed event and any related Common Base Events are linked by the same session ID.

To find and view related Common Base Events, use the following procedure.

Procedure

1. From within the administrative console, use the failed event manager to locate the failed event you want to investigate. See “Finding failed events” on page 38 for instructions on how to search for failed events.

- From the Failed Event Details page for that event, click **Browse Related Common Base Events**.

Results

The Common Base Event Browser opens in a new browser window and lists any Common Base Events related to the original failed event.

Deleting failed events

If you do not want to resubmit a failed event, or if you have failed events that have expired, use the failed event manager to delete them from the server. The failed event manager provides three options for deleting failed events.

Before you begin

You must be logged in as administrator or operator to perform this task.

About this task

To delete one or more failed events, use the following procedure.

Procedure

- Ensure that the failed event manager is open and that you have retrieved a list of the failed events on your system.
- From the failed event manager's Search Results page, do one of the following:
 - If you want to delete one or more specific failed events, select the check box next to each event and then click **Delete**.
 - If you want to delete only those failed events that have expired, click **Delete expired events**. Note that this deletes only the expired events in the current set of search results.
 - If you want to delete all failed events on the server, click **Clear all**.

Troubleshooting the failed event manager

This topic discusses problems that you can encounter while using the failed event manager.

Note: This topic does not discuss how to use the failed event manager to find, modify, resubmit, or delete failed events on the system. For information about managing failed events, see *Managing WebSphere Process Server failed events* in the information center.

Select the problem you are experiencing from the table below:

Problem	Refer to the following
I am having trouble entering values in the Search page's By Date tab	"Values in the By Date and From Date field automatically change to default if entered incorrectly" on page 21
I am having trouble deleting expired events	"Using the Delete Expired Events function appears to suspend the failed event manager" on page 22
I am having trouble with failed events not being created	"Failed events are not being created" on page 22

Values in the By Date and From Date field automatically change to default if entered incorrectly

The Search page's **From Date** and **To Date** fields require correctly formatted locale-dependent values. Any inconsistency in the value's format (for example, including four digits in the year instead of 2, or omitting the time) will cause the failed event manager to issue the following warning and substitute a default value in the field:

```
CWMAN0017E: The date entered could not be parsed correctly:  
your_incorrectly_formatted_date. Date: default_date is being used.
```

The default value of the **From Date** field is defined as January 1, 1970, 00:00:00 GMT.

Important: The actual default value shown in your failed event manager implementation will vary depending on your locale and time zone. For example, the From Date field defaults to 12/31/69 7:00 PM for a workstation with an en_US locale in the Eastern Standard Time (EST) time zone.

The default value for the **To Date** field is always the current date and time, formatted for your locale and time zone.

To avoid this problem, always enter your dates and times carefully, following the example provided above each field.

Using the Delete Expired Events function appears to suspend the failed event manager

If you use the Delete Expired Events button in situations where there are many failed events in the current search results, or where those events contain a large amount of business data, the failed event manager can appear to be suspended indefinitely.

In this situation, the failed event manager is not suspended: it is working through the large data set, and will refresh the results set as soon as the command completes.

Failed events are not being created

If the Recovery subsystem is not creating failed events, go through the following checklist of potential causes:

- Ensure that the wpsFEMgr application is running. If necessary, restart it.
- Ensure that the failed event manager's database has been created, and that the connection has been tested.
- Ensure that the necessary failed event destination has been created on the SCA system bus. There should be one failed event destination for each deployment target.
- Ensure that the Quality of Service (QoS) **Reliability** qualifier has been set to Assured for any Service Component Architecture (SCA) implementation, interface, or partner reference that participates in events you want the Recovery service to handle.

Recovering from a failure

Recovering from a failure requires an understanding of standard system processing in the event of a failure, as well as an understanding of how to analyze problems that may be the cause of a failure.

Overview of the recovery process

The recovery process encompasses a set of tasks that include both analysis and procedures.

When you must recover from a failure, these are the high-level steps to follow:

- Familiarize yourself with the possible kinds of failures. See *Triggers for recovery* for more information.
- Assess the state of the system. See *Assessing the state of the system* for more information.
- Form a hypothesis about what the problem is.
- Collect and analyze the data.
- Refer to other topics in this information center for instructions on fixing the problem.

Triggers for recovery

The need for solution recovery can result from a variety of triggers.

Situations from which solution recovery is necessary

Solution recovery is the process of returning the system to a state from which operation can be resumed. It encompasses a set of activities that address system failure or system instability that can be triggered by unforeseen circumstances.

You may need to perform solution recovery activities for the following circumstances:

- **Hardware failure**

Abnormal termination or system down can be caused by a power outage or catastrophic hardware failure. This can cause the system (all if not most JVMs) to stop.

In the case of a catastrophic hardware failure, the deployed solution may enter an inconsistent state on restart.

Hardware failures and environmental problems also account for unplanned downtime, although by far not as much as the other factors.

You can reduce the likelihood of hardware failures and environmental problems by using functions such as state-of-the-art LPAR capabilities with self-optimizing resource adjustments, Capacity on Demand (to avoid overloading of systems), and redundant hardware in the systems (to avoid single points of failure).

- **System is not responsive**

New requests continue to flow into the system but on the surface it appears that all processing has stopped.

- **System is unable to initiate new process instance**

The system is responsive and the database seems to work correctly. Unfortunately, new process instance creation is failing.

- **Database, Network or Infrastructure Failure**

In the case of fundamental infrastructure failure, the solution may require administration to restart/resubmit business transactions after the infrastructure failure is resolved.

- **Poor tuning or a lack of capacity planning**

System is functional but is severely overloaded. Transaction time-outs are reported and there is evidence of an overflow of the planned capacity.

Incomplete capacity planning or performance tuning can cause this type of solution instability.

- **Defects in application module development**

The modules that are part of a custom developed solution can have bugs. These bugs can result in solution instability and failed services.

Bugs in a custom developed solution can result from a variety of situations, including (but not limited to) the following:

- Business data that was not planned for or unforeseen by the application design.
- An incomplete error handling strategy for the application design.

A detailed error handling design can reduce solution instability.

- **WebSphere software defect**

A defect in the WebSphere product causes a backlog of events to process or clear.

Assessing the state of the system

The first thing to do when an abnormal condition occurs is to take the *pulse* of the overall system and get a feel for how much or how little of the system is operational and how much of it is rendered 'out of service' by whatever the external stimuli was that caused this condition.

Address a predefined set of questions to assess the extent of the outage. The following list provides examples of predefined questions designed to help you gather the appropriate information:

1. Is this system still performing work?

Determine if system is still operational. Often times, a system can be operational, but because of overload or inappropriate tuning, or both, the system is not completing tasks quickly and/or is attempting to do work that is indeed failing.

The litmus test for each of these questions will be specific to the nature of the deployed solution.

2. What special error handling support is built-in to the application?

If there is a lot of automated retry and various support logic, then the application itself might shield some errors from manifesting the IT operator.

These conditions must be known and documented for reference by the recovery team.

Things you can do to help assess the state of the system include the following:

1. Check to see if the server is at least running.

Do you see the PID or get a positive feedback from the deployment manager via the administrative console?

2. Check to see if there are locks in the database(s) or any unusual database traffic.

Most databases will have facilities to look at locks. Depending on the deployment topology, there may be multiple databases.

- Messaging Engine Database
- Business Process Container Database
- WebSphere Process Server Common Database (Failed Events and Relationship data)

3. Check to see what the status of the messaging system is.

Check for events or messages in the following locations:

- Business Process Choreographer Hold and Retention Destinations
- Number of failed events
- Number of messages on the solutions module destinations

4. Check to see if the database is functioning.

Can you perform some simple SELECT operation, on unlocked data in a reasonable amount of time?

5. Check to see if there are errors in the database log.


If the database is not working properly, then recovering the database (so that it can at least release locks and perform simple selects) is vital to system recovery.

If the messaging system is not working properly, then recovering the messaging subsystem, so that it can at least be viewed and managed, is also vital to system recovery.

Note: A 'bottoms up' approach is not always conclusive. However, chances of successful recovery vary based on these basic activities.

From these basic procedures and health check kinds of activities, start to look at some specific situations. Patterns will be described, specifics will be given and insights as to what is going on under the covers will be provided.

Realize that this situational analysis is a read-only activity. While it provides vital information from which to determine the appropriate recovery actions, it should not change the state of the system under review. It is impossible to predict and provide prescriptive actions for all possible causes of a system outage. For example, consider the following decision tree:



There are many broad categories to investigate in the event of an unplanned outage. These broad categories will have sub categories and so on. Defining prescriptive actions for each node and the subsequent node will depend on the results of each investigation. Because this type of relationship is difficult to convey in a document form, using a support tool such as *IBM Guided Activity Assist* to interactively walk you through the investigative and decision making process is recommended. As you progress from the top to each child node, it is important to conduct the appropriate level of situational analysis.

Recovery: Analyzing the problem

For all unplanned system events, a set of basic recovery procedures can be leveraged at the point of identification.

There are several well defined steps to situational analysis. The steps are listed below.

1. Define the question
2. Gather information and resources (observe)
3. Form hypothesis
4. Perform experiment and collect data
5. Analyze data
6. Interpret data and draw conclusions that serve as a starting point for new hypothesis

For each production scenario the symptoms that initiate a recovery action may vary.

It is important to follow the guidelines for situational analysis and take the corrective action relative to the symptoms that are presented.

Situational analysis

Situational analysis is the cyclical execution of the scientific method and can take into account various situations that will initiate a recovery procedure.

The following list is of the different types of situations that will initiate a recovery procedure:

- Abnormal termination or system down
A power outage or catastrophic hardware failure has caused the system (all if not most JVMs) to stop.
- System is not responsive
New requests continue to flow into the system but on the surface it appears that all processing has stopped.
- System is functional but is severely overloaded
Transaction time-outs are reported and there is evidence of an overflow of the planned capacity.
- System is unable to initiate new process instance
The system is responsive and the database seems to work correctly. Unfortunately, new process instance creation is failing.

Recovery: First steps

Administrators can facilitate solution recovery processes by following a first steps checklist of general practices.

The following list describes actions that you **SHOULD NOT TAKE** under normal circumstances when trying to recover a solution.

Note: There could be special situations for which you might need to perform some of the actions listed below. However, you should never initiate any of these actions without first consulting with the WebSphere Process Server support organization.

- Do not delete the transaction log file
The transaction (tranlog) log file stores critical transactional data that is written to databases. It is an internal file that WebSphere Application Server uses to manage in-flight transactions and attempt to recover them should the server lock up.
- Do not have the transaction logs local on the cluster members
Put the transaction logs on a shared drive. This is the only way to allow peer recovery, which helps minimize the downtime during recovery.
- Do not attempt database operations where the result set is large enough to create additional resource contention (OutOfMemory)
- Avoid performing Business Process Choreographer Explorer operations that return large result sets.
- Avoid executing administrative scripts on process instances without considering the result set size.
- Do not drop and/or re-create databases in production
- Do not uninstall applications as part of your standard recovery procedures
You should only uninstall applications with the direction from the IBM support organization.

- Do not enable too much trace if the system is overloaded.
Too much trace will cause a slowdown in system throughput and might cause transaction time-outs. Too much trace can often add to the problems that need to be addressed, rather than providing insight as to how to solve the original problems.
Get immediate assistance from IBM support to define the correct trace specification.
- Do not experiment or try out new scripts or new commands on production systems.
- Do not run your production servers in *development mode*
Enabling the **Run in development mode** option may reduce the startup time of an application server. This may include JVM settings such as disabling bytecode verification and reducing JIT compilation costs.



- The following list describes the recommended actions when it comes to recovery.
- Always take a *snapshot* of the configuration tree, the PI file of the application in question and the logs which are available.
Logs may be overwriting themselves depending on the configuration. Capturing a set early and often is an important step for postmortem analysis. See the topic on *IBM Support Assistant (ISA)* for details on the IBM Support Assistant, which helps with this type of activity.
 - Always understand your database settings, especially related to database transaction log file size, connection pools, and lock timeouts.

Failed-event locations: Where does the data go?

For all (production and test) recovery activities there are a finite number of locations in the solution where events accumulate.

By adhering to guidelines and preventive measures described in *Planning error prevention and recovery*, all business events and associated data will reliably accumulate in one of these locations.

If you do not adhere to sound architectural and application development practices, then a percentage of inflight events may end up in an inconsistent state, from which recovery cannot be attained. Under such circumstances, (presumably

identified during testing cycles) post recovery investigation and clean up is necessary to correct the issue so that future recovery activities are completely successful.

In order to accurately describe the following scenarios, it is important to put the information in the context of a use case.

Use case: recovering data from failed events

A use case provides a context for a recovery scenario. In the use case, a business has an application that receives a request to create a new Account.

The solution is comprised of multiple modules as recommended through module best practices.

The first module mediates the request and delegates work to an Account Creation process. In the example below we have implemented the solution as separate modules where the request is passed between the mediation module (AccountRouting) and the processing module (AccountCreation) via an SCA import/export. See the following screen capture for an illustration of the two modules.

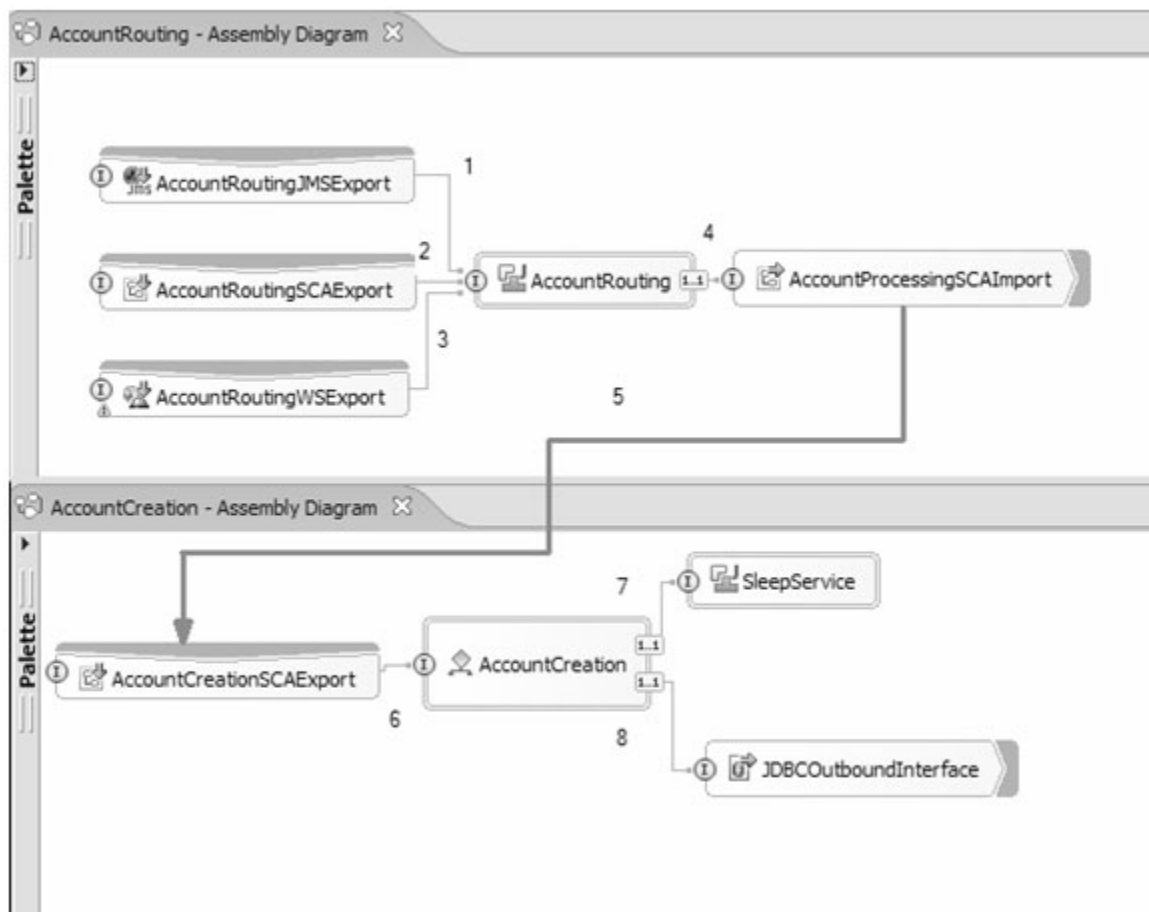


Figure 1. Assembly diagram of account routing process

From the assembly diagram shown in Figure 1, you can begin to see at what locations in the flow that failures might occur. Any of the invocation points in the

assembly diagram can propagate or involve a transaction. There are a few areas in the flow where data will collect as a result of application or system failures.

In general, transaction boundaries are created and managed by the interaction (synchronous and asynchronous) between components and import/export bindings and their associated qualifiers. Business data accumulates in specific recovery locations most often due to transaction failure, deadlock or rollback.

Transaction capabilities within WebSphere Application Server help WebSphere Process Server enlist transactions with service providers. These enlisted interactions are particularly important to understand with respect to import and export bindings. Understanding how imports and exports are used within your specific business cases is important in determining where events in need of recovery accumulate.

An error handling strategy should define interaction patterns, transactions used, import and export usage before developing the application. The solution architect should identify the preferences to use, guidelines, that are then used as the application is created. For example, the architect needs to understand when to use synchronous vs. asynchronous calls, when to use BPEL fault handling and so forth. The architect needs to know whether or not all services can participate in transactions, and for those services that can not participate, how to handle compensation if problems are encountered.

Additionally, the application shown in the assembly diagram in Figure 1 on page 61 leverages connectivity groups and module development best practices. By leveraging this pattern we now have the ability to stop the inbound flow of new events by stopping the AccountRouting module.

The following sections address the location of business data in the case of failure and recovery.

Business Flow Manager or Human Task Manager

In our business case, we leverage a BPEL process for AccountCreation process.

With regard to recovery, there are a some questions you need to ask yourself with respect to BPEL and Human Task management as follows:

1. What type of process is being run (short running or long running, business state machine, human task) ?
Short running processes are known as microflows.
2. Is the process developed properly and using fault handling to promote data integrity?
3. How are the invocation patterns and unit of work properties configured to predict and control transaction boundaries?

Knowing the answers to these questions will impact your recovery strategy for invocations 7 and 8 shown in the assembly diagram, as highlighted in the screen capture below:




Figure 2. Assembly diagram of account routing - invocations 7 and 8

Stateful components, such as Long Running BPEL processes and Business State Machines, involve many database transactions where process activity changes and state changes are committed to the database. The work progresses by updating the database and placing a message on an internal queue that describes what is to be done next. More information on Macro flow transactions is available in the information center topic titled *Transactional behavior of long-running processes*.

If there are problems processing messages that are internal to the Business Flow Manager, these messages are moved to a *Retention Queue*. The system attempts to continue to process messages. If a subsequent message is successfully processed, the messages on the Retention Queue are resubmitted for processing. If the same message is placed on the Retention Queue five times, it is then placed on the Hold Queue. Information, such as what internal queues are used or retry algorithms for these queues is described in detail in the information center topic titled *Recovery from infrastructure failures*.

Additional information about viewing the number of messages and replaying messages can be found in *Replaying Messages from the Retention Queue / Hold Queue*.

Failed event manager

The Failed event manager (FEM) is used to replay events or service invocation requests that are made asynchronously between *most* component types.

Failed events are created if the AccountRouting component makes an asynchronous call to the SCA Import binding AccountCreationSCAImport and a ServiceRuntimeException is returned.

It is important to note that failed events are not generated in most cases where BPEL is the client in the service interaction. This means that the invocation for 7 and 8 (as shown in Figure 2 on page 63) will not typically result in a failed event. BPEL provides fault handlers and other ways to model for failure. For this reason, if there is a `ServiceRuntimeException` (SRE) failure calling "JDBCOutboundInterface", the SRE is returned to the BPEL for processing. The error handling strategy for the project should define how runtime exceptions are consistently handled in BPEL.

It is important to note however, failed events are created for asynchronous response message for the BPEL client if these messages can not be delivered to the process instance due to an infrastructure failure.

The following diagram illustrates how the Failed event manager component works. Descriptions of the processing associated with each numbered step are provided following the diagram.




Figure 3. Failed event manager processing

Failed event manager processing

1. The source component makes a call using an asynchronous invocation pattern
2. The SCA MDB picks the message up off the SCA destination
3. The SCA MDB makes the call to the correct target component
4. The target component throws a `ServiceRuntimeException`
5. The SCA MDB transaction rolls back to the SCA destination
6. The exception information is stored to the Failed event manager database with a status of *not confirmed*
7. The invocation is retried by the SIBus n number of times

The retry limit default value is 5 - one original and 4 retries. You can change the default value in the administrative console. For example, given an SCA module M, you could navigate to **Buses** → **SCA.SYSTEM.<CELL>.BUS** → **Destinations** → **sca/M** and change the value in the *Maximum failed deliveries* field.

8. After the number of retries reaches the specified limit, the message is moved to the FEM destination.
9. The Failed Event Manager database picks up the message
10. The Failed Event Manager database updates the failed event in the database and the status is set to *failed*.

When are "failed events" created?

As stated, failed events are neither created for synchronous invocations nor typically for two-way business process interactions.

Failed events are generally created when clients use an asynchronous invocation pattern and a `ServiceRuntimeException` is thrown by the service provider.

If everything is done synchronously and in the same transaction, data is not collected anywhere. Instead it is all rolled back to the client that made the call. Where ever a commit is occurs, data collects. If the calls are all synchronous, but there are multiple commits, then these commits become an issue.

In general, you should use asynchronous processing calls or long running BPEL if multiple transactions are needed. So each ASYNC call is a chance for data to collect. Long running BPEL process are a collection point.

*Table 6. Invocation patterns and relationship to the creation of failed events: **Service Business Exceptions***

Invocation Pattern	Failed Event Created Y/N?	Notes
Synchronous	No	Failed events are not created for service business exceptions or when using a synchronous pattern
Asynchronous - One Way	No	By definition, one-way invocations cannot declare faults, meaning, it is impossible to throw a <code>ServiceBusinessException</code> .
Asynchronous - Deferred Response	No	Failed Events are not created for service business exceptions
Asynchronous - Callback	No	Failed Events are not created for service business exceptions

*Table 7. Invocation patterns and relationship to the creation of failed events: **Service Runtime Exceptions***

Invocation Pattern	Failed Event Created Y/N?	Notes
Synchronous	No	Failed events are not created for service runtime exceptions or when using a synchronous pattern.
Asynchronous - One Way	Yes	
Asynchronous - Deferred Response	Yes	

Table 7. Invocation patterns and relationship to the creation of failed events: **Service Runtime Exceptions** (continued)

Invocation Pattern	Failed Event Created Y/N?	Notes
Asynchronous - Callback	Yes	
BPEL - Two Way	No	Failed Events are not created when the source component is a business process. Note: For an asynchronous call, if the response can not be returned to BPEL, then a failed event is created.
BPEL - One Way	Yes	

For additional information, review the information center topic titled *Managing failed events*.

Additional information about viewing and resubmitting failed events can be found in section *Resubmitting Failed Events*.

Service Integration Bus Destinations

Messages that are waiting to be processed may accumulate in a few Service Integration Bus (SIBus) destinations. For the most part these destinations are "system" destinations. Messages within these destinations typically are a mixture of three types as follows:

- Asynchronous requests for processing
- Asynchronous replies to requests
- Asynchronous messages that failed deserialization or function selector resolution

Note: Asynchronous replies can be valid Business Objects or faults returned as a result of a request.

SCA Module Destination

Again, referring back to our business case.

There would be two "SCA Module" destinations in the solution:

- sca/AccountRouting
- sca/AccountCreation

These destinations are created when the module is deployed to an application server or a cluster.

There are rare opportunities for messages to accumulate in these destinations. The accumulation of messages in these locations is a strong indication that there maybe a performance problem or an application defect. Investigate immediately. It is important to monitor the depth of the module destinations (with your chosen IT monitoring solution) as a back up of messages could lead to a system outage or a prolonged recycle time.

We call these "SCA Module" destinations because the generated name is the same as the module name with the additional "sca/". These destinations are pivotal in the functioning of SCA asynchronous invocations (brokering requests and responses). There are a varying number of additional destinations that are

generated during application installation on the SCA.SYSTEM bus but for the purpose of the discussion we'll be addressing the importance of the "SCA Module" destination.

System Integration Bus Retry

As we learned above, the FEM has a built-in retry mechanism with the SCA message driven bean (MDB). This retry behavior can be controlled by modifying the "Maximum Failed Deliveries" attribute on the module destination.

Note: Typically, there is no reason to adjust this retry capability. The information provided here is for completeness.

Referring to our business case, there are a number of SI Bus destinations created by SCA to support asynchronous communication.

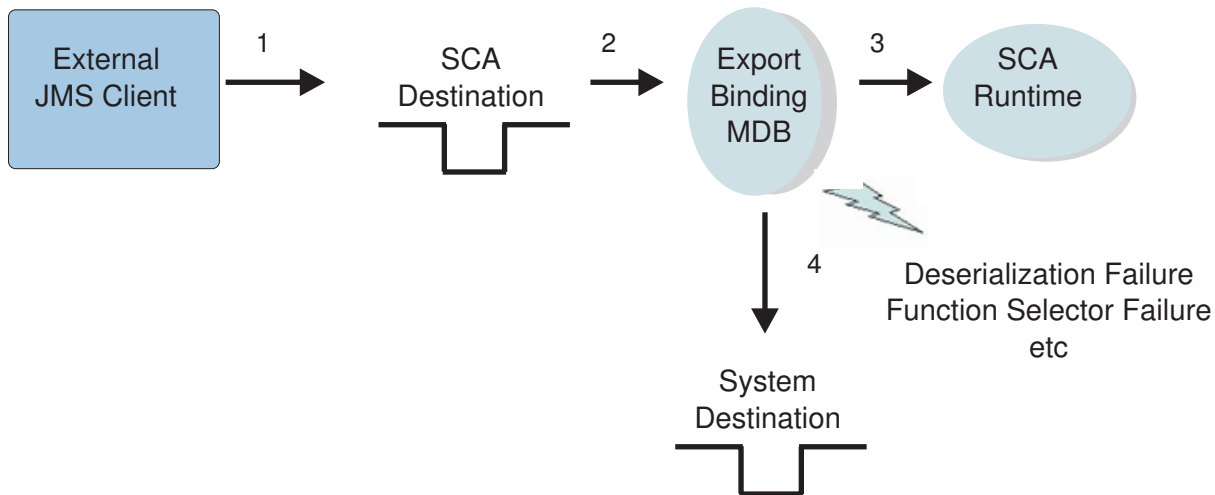
As we have learned, one of these destinations is called "sca/AccountRouting" You can adjust the number of retries that happen during a ServiceRuntimeException of an asynchronous service invocation by changing the value of the "Maximum Failed Deliveries" property via the admin console. However, you may not set the value less than 2 in modules with a BPEL process. The second delivery is required to return ServiceRuntimeExceptions back to the BPEL for processing.

System Exception Destinations

The failed event manager is one place where we can look to administer failures. When dealing with Imports and Exports that are JMS or EIS based, we have to consider another important location.

Destinations on the SCA.Application bus are configured to route failed messages to the SIB system exception destination for that bus. Thus, if a JMS export picks up a message from the SCA Application bus and runs into a rollback situation, the failed message will be routed to the SIB system exception destination instead of the WBI recovery exception destination. This scenario differs from the failed event discussion above in that a failure to deserialize a message on the SCA.Application bus will not result in a failed event. There is a system exception destination on every bus within the solution. These destinations must be monitored and administered much like the "dead letter queue" common to MQ infrastructures.

Consider the following scenario.



An external JMS client places a message on an inbound queue exposed via a JMS Export. The JMS Export Binding MDB picks up the message for processing. From here, one of two things will happen:

1. The JMS export successfully parses the message and determines which operation on the interface to invoke at which point the message is sent to the SCA runtime for processing.
2. The JMS export fails to recognize the message body as a valid business object or the JMS export binding *deserializes* the message body but is unable to determine the appropriate operation on the interface to invoke. At this point the message is placed on the System Exception Destination for the bus.

We can have this type of failure when trying to receive requests from the AccountRoutingJMSEExport (1). This export is a JMS export and there is a possibility that events can accumulate on the System Exception Destination on the SCA.Application.Bus. Use the chosen IT monitoring solution to observe the depth of this destination.

Failed Event Manager and SIB Destinations

For WebSphere Process Server, the exception destination is set to the WebSphere Process Server exception destination queue. This queue follows a naming convention as follows:

```

Node name: WPSNode
Server name: server1
Recovery exception destination: WBI.FailedEvent.WPSNode.server1
  
```

In general, all the destinations created on the SCA.System bus will be configured to route failed messages to the recovery exception destination.

When a system failure occurs, in addition to capturing the failed message in this exception destination, the WebSphere Process Server recovery feature also generates a failed event that represents the system error and stores it into the Recovery database as described in the Failed Event Manager section of this document.

Summary

In summary, WebSphere Process Server provides administrative capabilities above and beyond the underlying WebSphere Application Server platform. Proper

measures should be made to understand and use these capabilities along with following the guidance provided in the Planning error prevention section of *Planning error prevention and recovery*.

Table 8. Administrative capabilities to help manage failures

Administrative Capability	Bundled With WebSphere Process Server Y/N?	Summary
Business Process Choreographer Explorer	Yes	Read/Write/Edit/Delete Access. This is the central place to administer business processes and human tasks.
Failed Event Manager	Yes	Read/Edit/Delete Access. This is the central place to administer Service Runtime Exceptions and other forms of infrastructure failures.
Service Integration Bus Browser	Yes	Read/Delete. Use the Service Integration Bus Browser on the administrative console for browsing and performing day-to-day operational tasks on service integration buses.

Note: The number of events or records that can be simultaneously administered by these tools are specific to external factors such as memory allocation, result sets and DB tuning, connection timeout. Run tests and set the appropriate thresholds to avoid exceptions (OOM, TransactionTimeout).

Related concepts

“Retention queues and hold queues” on page 76

When a problem occurs while processing a message, it is moved to the retention queue or hold queue.

Recovery troubleshooting tips

This section provides a list of tips for troubleshooting the recovery process.

Restarting deployment environments

As one step in a recovery process, you may need to restart of you deployment environment.

About restarting deployment environments

The procedure to restart a deployment environment varies depending on the topology. Topologies are based on system configuration patterns, each pattern designed to meet particular business requirements.

WebSphere Process Server supports a set of predetermined deployment environment configuration patterns. If none of the patterns meet your requirements, you can plan and create your own customized deployment environment.

In any given deployment environment configuration pattern there are a number of servers running as JVM processes. In general there are three types of servers as follows:

- Messaging Servers

Messaging servers are responsible for providing the Service Integration Bus (SIB) messaging infrastructure.

- **WebSphere ESB Servers**
Servers with profiles capable of only hosting and running mediation modules.
- **WebSphere Process Servers**
Servers with profiles capable of hosting and running all module types. This profile hosts the Business Process Choreographer component.
- **Support Servers**
This server is responsible for providing support and monitoring services such as the Common Event Infrastructure CEI.

The deployment patterns differ in how you group and organize all the functional components, so that the pattern can address your business requirements in the most cost effective fashion. For more advanced and highly available environments, the servers would reside in clusters that are distributed across physical resources.

General practice for restarting servers as part of a recovery operation

A general model for starting servers is to start the messaging servers first, then the support servers and lastly the WebSphere Process Server servers. Each application architecture may have specific dependencies between application components that need to be taken into consideration.

Shutdown basically happens inverse to the startup procedure, starting with the application server clusters and ending with shutting down the messaging infrastructure after it has had time to quiesce and process any inflight transactions.

Related tasks

 [Choosing your deployment environment pattern](#)

You can configure your deployment environment by choosing one of the IBM-supplied topology patterns or by creating your own custom deployment environment. This topic section lists and describes the available IBM-supplied topology patterns and presents considerations for choosing a topology.

 [Planning your deployment environment](#)

Setting up your deployment environment involves many decisions that affect everything from the number of physical servers to the type of pattern you choose. Each decision will affect how you set up your deployment environment.

Related information

 [WebSphere Process Server and WebSphere Enterprise Service Bus deployment patterns: Selecting your deployment pattern](#)

Viewing the service integration bus

Use the Service Integration Bus browser on the administrative console to view the service integration bus.

Before you begin

Make sure you understand how the Service Component Architecture (SCA) system bus is used.

About this task

The Service Integration Bus Browser provides a single location for browsing and performing day-to-day operational tasks on service integration buses.

Viewing the service integration bus is a useful way to determine if messages are accumulating on the SCA Module destinations.

The accumulation of messages on the SCA Module destinations is a strong indication that there maybe a performance problem or an application defect.

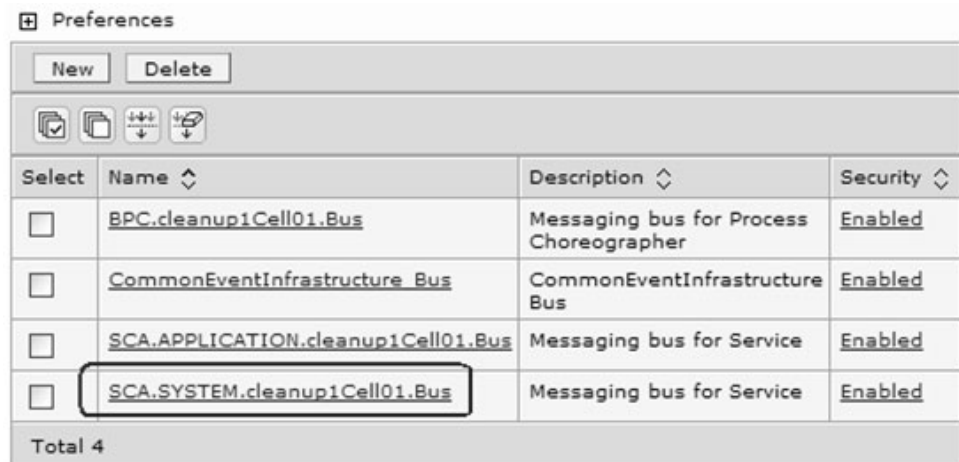
It is a good idea to periodically view the messages and determine if there are any messages have become locked for an extended duration of time as this may indicate that there are "indoubt transactions".

Procedure

1. From the administrative console, expand **Service integration**.
2. Select **Buses**.



3. Select the appropriate messaging bus for the service. The following example shows the messaging bus named `SCA.System.cleanup1cell101.bus` highlighted, where `cleanup1cell101` is the name of the cell.

A screenshot of a 'Preferences' dialog box. It features a table with columns for 'Select', 'Name', 'Description', and 'Security'. The table lists four messaging buses. The last row, 'SCA.SYSTEM.cleanup1Cell01.Bus', is highlighted with a red rectangle. The table also includes 'New' and 'Delete' buttons at the top, and a 'Total 4' summary at the bottom.

Select	Name	Description	Security
<input type="checkbox"/>	<u>BPC.cleanup1Cell01.Bus</u>	Messaging bus for Process Choreographer	<u>Enabled</u>
<input type="checkbox"/>	<u>CommonEventInfrastructure Bus</u>	CommonEventInfrastructure Bus	<u>Enabled</u>
<input type="checkbox"/>	<u>SCA.APPLICATION.cleanup1Cell01.Bus</u>	Messaging bus for Service	<u>Enabled</u>
<input type="checkbox"/>	<u>SCA.SYSTEM.cleanup1Cell01.Bus</u>	Messaging bus for Service	<u>Enabled</u>

4. Select **Destinations**



- Review the relevant information. You should look at the destinations named `sca/XYZ`, where `XYZ` is the name of the module. For example, for modules named `AccountRouting` and `AccountCreation`, you would look for the following destinations:

<input type="checkbox"/>	<u>sca/AccountCreation</u>
<input type="checkbox"/>	<u>sca/AccountCreation/component/AccountCreation</u>
<input type="checkbox"/>	<u>sca/AccountCreation/component/SleepService</u>
<input type="checkbox"/>	<u>sca/AccountCreation/export/AccountCreationSCAExport</u>
<input type="checkbox"/>	<u>sca/AccountCreation/exportlink/AccountCreationSCAExport</u>
<input type="checkbox"/>	<u>sca/AccountCreation/import/JDBCOutboundInterface</u>
<input type="checkbox"/>	<u>sca/AccountCreation/import/sca/dynamic/import/scaimport</u>
<input type="checkbox"/>	<u>sca/AccountCreation/import/sca/dynamic/import/vsimport</u>
<input type="checkbox"/>	<u>sca/AccountRouting</u>

- Select the link text for the destination that you are interested in viewing. This will link you to a general properties page for the destination that you want to view.
- From the general properties page of the destination, select the **Queue points**

Configuration

General Properties

Identifier
sca/AccountCreation

UUID
5D2AB86F4EDEC81E01F34714

Type
Queue

Description
active:sca/AccountCreation

Message points

- Queue points
- Mediation points

Additional Properties

- Context properties
- Mediation execution points

- From the Queue points page, select the link for the message point.

Buses > SCA.SYSTEM.cleanup1Cell01.Bus > Destinations > sca/AccountCreation > Queue points

The message point for a queue, for point-to-point messaging.

Preferences

Identifier

sca/AccountCreation@default.Messagingq.000-SCA.SYSTEM.cleanup1Cell01.Bus

Total 1

- Select the **Runtime** tab.
From this screen you can see the current message “depth” and the threshold.
Selecting the **Messages** link will allow you to view the message contents.

The message point for a queue, for point-to-point messaging.

Configuration Runtime

Refresh

General Properties

Identifier
sca/AccountCreation

Run-time ID
5D2AB86F4EDEC81E01F34714_QUEUE_28000008

High message threshold
50000

Send allowed

Current message depth
0

OK

Additional Properties

- Messages
- Known remote queue points

Ideally, use an appropriate IT monitoring tool and set alert thresholds for these destinations. The threshold value would be established during the performance test phase for the application.

Messages on a production system should never be deleted unless explicitly directed to do so by the SCA L3 team.

Related concepts

SCA system bus

The *SCA system bus* is a service integration bus that is used to host queue destinations for Service Component Architecture (SCA) modules. The SCA run time, which supports mediation modules, uses queue destinations on the system bus as an infrastructure to support asynchronous interactions between components and modules.

Service integration buses for WebSphere Process Server

A service integration bus is a managed communication mechanism that supports service integration through synchronous and asynchronous messaging. A bus consists of interconnecting messaging engines that manage bus resources. It is one of the WebSphere Application Server technologies on which WebSphere Process Server is based.

Service Integration Bus Browser

The Service Integration Bus Browser provides a single location for browsing and performing day-to-day operational tasks on service integration buses.

Related tasks

“Resolving indoubt transactions” on page 79

Transactions can become stuck in the indoubt state indefinitely due to exceptional circumstances, such as the removal of a node causing messaging engines to be destroyed.

Related information

SCA resources

Considerations for Service Component Architecture support in servers and clusters

Servers and clusters can support Service Component Architecture (SCA) applications, application destinations, or both.

Capturing javacore

There are a number of methods that you can use to capture a javacore from an IBM JDK and thread dumps for non-IBM JDKs.

Capturing javacore

A javacore dump, or a thread dump as it is also called, is one of the primary problem determination documents that an application server creates.

1. Use wsadmin to produce a javacore in the Profile directory:

- a. For Windows®:

```
<PROFILE_DIR>\bin\wsadmin.bat [-host host_name] [-port port_number]
[-user userid -password password] -c
"$AdminControl invoke [$AdminControl queryNames WebSphere:name=JVM,process=server1,*]
dumpThreads"
```

- b. For Unix (IBM JDKs):

```
<PROFILE_DIR>>/bin/wsadmin.sh[-host host_name]
[-port port number] [-user userid -password password] -c
"\$AdminControl invoke [\$AdminControl queryNames WebSphere:name=JVM,process=server1,*]
dumpThreads"
```

Note: The braces, [] around the AdminControl queryNames command are part of the command, and not used to signify optional parameters as is the case for braces around host, port and user. The process name: server1 may need to be change to fit your configuration.

2. A signal can be sent to the server process:

a. Windows:

A launch script must be used to start the server process to allow the signal to be passed to the process. This does require special setup before starting the server.

1) <PROFILE_DIR>\bin\startServer.bat server1 -script SERVER1.bat

2) b. SERVER1.bat

The server process will start in a command window. You will need to check the logs to verify that the server has successfully started since the intermediate JVM process which usually starts the server process is not used.

3) <CTRL><BREAK>

Issue a <CTRL><BREAK> into the command window where the server process is running. A javacore will be produced.

b. **Unix (all JDKs):** kill -3 <pid>

Where <pid> is the process id of the WebSphere Process Server. For IBM JDKs a javacore will be produced in the <PROFILE_DIR>directory.

For non-IBM JDKs, a thread dump will be written to the native_stdout.log.

3. An alternative method to dumping a windows core file is to use jvmdump.

This does not require special setup before starting the server. However, it does require a special executable from the JVM team. The jvmdump.exe program can be requested by sending a note to jvmcookbook@uk.ibm.com. The advantage of this method is additional information can be obtained about native code being executed within JVM. The format of the dump differs from the IBM javacores.

- jvmdump.exe <PID>

- <WAS_HOME>\java\jre\bin\jextract.exe <core.name.dmp>

- <WAS_HOME>\java\jre\bin\jdumpview.exe

- set dump <core.name.dmp>.zip

- display thread


Displays the current executing thread at the time of the dump

- c. display thread *

Display all of the threads from the dump.

For more details about the jdumpview utility consult the Diagnostics Guide for the IBM Developer Kit and Runtime Environment, Java Technology Edition, Version 5.0.

Related information

 [Configuring the hang detection policy](#)

Servers and recovery mode processing

When you restart an application server instance with active transactions after a failure, the transaction service uses recovery logs to complete the recovery process.

These recovery logs, which each transactional resource maintains, are used to rerun any Indoubt transactions and return the overall system to a self-consistent state. An *indoubt transaction* is one that has encountered environmental or other errors

during commit processing. Logging occurs for normal inflight transactions, but those log entries are removed upon successful commit processing.

This recovery process begins as soon as all of the necessary subsystems within the application server are available during a server startup. If the application server is not restarted in recovery mode, the application server can start accepting new work as soon as the server is ready, which might occur before the recovery work has completed. This might be okay in many cases, but the more conservative option is provided here. To be clear, recovery will run on a server restart even if the server is started in 'normal' start model.

For information on how to start a server, see the Starting a server topic in the WebSphere Process Server information center.

Related concepts

“Profile-specific log files” on page 19

There are log files detailing the characteristics and runtime activities of individual profiles. These log files are located within the logs directory of the profile path.

Related tasks



Starting managed servers

Start the application server process in order to run applications on the managed server.

Retention queues and hold queues

When a problem occurs while processing a message, it is moved to the retention queue or hold queue.

You can perform administrative actions on the messages in the retention queue and hold queue using either the administrative console or through scripting.

In some cases, viewing and replaying messages on the retention queue or the hold queue can be part of a recovery procedure.

Related concepts





“Use case: recovering data from failed events” on page 61

A use case provides a context for a recovery scenario. In the use case, a business has an application that receives a request to create a new Account.

“Managing failed events” on page 33

The WebSphere Process Server Recovery service captures data about failed events. You can then use the failed event manager to view, modify, resubmit, or delete the failed event.

Related information

-  Business processes: Recovery from infrastructure failures
-  Failed event manager console help field descriptions
-  Querying and replaying failed messages, using the administrative console
-  Querying and replaying failed messages, using administrative scripts

Business Process Choreographer maintenance and recovery scripts

There are several maintenance-related scripts for Business Process Choreographer. Run these maintenance scripts as part of a general maintenance policy to help maintain database performance, or as part of a recovery process as deemed necessary.

You should run these scripts to remove from the database the templates and their associated objects, as well as completed process instances, that are not contained in any corresponding valid application in the WebSphere configuration repository.

There is also the possibility of having invalid process templates. This situation can occur if an application installation was canceled or not stored in the configuration repository by the user.

WebSphere Process Server also provides a service that automates Business Process Choreographer cleanup. You can run that service from the administrative console.

Use the following scripts for Business Process Choreographer recovery maintenance:

- `deleteInvalidProcessTemplate.py`

Run this script to delete, from the Business Process Choreographer database, business process templates that are no longer valid.

Note: These templates usually have no impact. They are not shown in Business Process Choreographer Explorer.

You cannot use this script to remove templates of valid applications from the database. This condition is checked and a `ConfigurationError` exception is thrown if the corresponding application is valid.

- `deleteInvalidTaskTemplate.py`

Run this script to delete, from the Business Process Choreographer database, human task templates that are no longer valid.

You cannot use this script to remove templates of valid applications from the database. This condition is checked and a `ConfigurationError` exception is thrown if the corresponding application is valid.

- `deleteCompletedProcessInstances.py`

Run this script when all completed process instances have to be deleted.

A top-level process instance is considered completed when it is in one of the following end states:

- Finished
- Terminated
- End
- Failed

You can specify the criteria to selectively delete top-level process instances and all their associated data (such as activity instances, child process instances, and inline task instances) from the database.

Note: When running these scripts from the command line, make sure the SOAP client timeout is set high enough to complete the requested operation for the WAS Admin client.

Deleting an allotment of completed process instances

You can delete an allotment of process instances from the development environment.

Using a script that wrappers the provided `deleteCompletedProcessInstances.py`

By editing and placing correct user names, passwords, and paths in this wrapper script, you can delete an allotment of process instances from the development environment.

Carefully selecting an adequate time slice prevents SOAP timeout exceptions when communicating with the deployment manager.

The “adequate time slice” of administrable instances depends on many factors including, but not limited to, the following:

- JVM tuning and memory allocations
- Transaction log configuration for the database server
- SOAP connection Time-Out configuration

Example

For example, after altering the script and running the command as:







```
wsadmin.<bat|sh> -user<USERNAME> -password<PASSWORD> -f loopDeleteProcessInstances.py  
2008-04-02T21:00:00 3600
```

This command will run `deleteCompletedProcessInstances.py` while increasing the completed before time stamp by one hour (60 minutes * 60 seconds) after every execution.

The `deleteCompletedProcessInstances.py` script has a time stamp parameter which can be used to control the number of instances being deleted. The smaller the interval, the fewer instances will be deleted per invocation of the `deleteCompletedProcessInstances.py`. This can be useful in situations where deletion of multiple process instances encounter transaction timeouts. The most common cause for transaction timeouts during process deletion involve the following:

- an untuned database
- an overloaded system
- attempting to delete “too many” process instances at once

Related information

-  [Process instances](#)
-  [Using scripts to administer Business Process Choreographer](#)
-  [Deleting process templates that are unused](#)
-  [Deleting completed process instances](#)
-  [Deleting human task templates that are unused](#)
-  [Configuring the cleanup service and cleanup jobs](#)

Resolving indoubt transactions

Transactions can become stuck in the indoubt state indefinitely due to exceptional circumstances, such as the removal of a node causing messaging engines to be destroyed.

Before you begin

Use the procedure to resolve indoubt transactions only if you have tried other procedures (such as restarting the server in recovery mode), unsuccessfully.

About this task

When a transaction is stuck in the indoubt state, it must either be committed or rolled back so that normal processing by the affected messaging engine can continue.

You can use the administrative console to display the messages causing the problem by Listing messages on a message point.

If there are messages related to an indoubt transaction, the identity of the transaction displays in a panel associated with the message. You can then resolve the transaction in one of the following ways:

- Using the server's transaction management panels
- Using methods on the messaging engine's MBean

You should first attempt to resolve the indoubt transaction using the application server transaction management panels. If this does not work, then use methods on the messaging engine's MBean. Both procedures are described below.

Procedure

1. **Using the application server transaction management panels to resolve indoubt transactions**
 - a. Navigate to the transaction management panels in the administrative console

Click **Servers** → **Application servers** → [Content Pane] → *server-name* → [Container Settings] Container Services → Transaction Service → Runtime → Imported prepared transactions - Review

- b. If the transaction identity appears in the resulting panel, you can commit or roll back the transaction

Choose the option to roll back the transaction

If the transaction identity does not appear in the panel, the transaction identity was not enlisted with the Transaction Service on the server. In this case only, you should use methods on the MBean (as described in the next step) to display a list of the identities of the indoubt transactions managed directly by the messaging engine.

2. Using methods on the messaging engine's MBean to resolve indoubt transactions

CAUTION:

Only perform this step if you were unable to display the transaction identity by using the server's transaction management panels

- a. The following methods on the messaging engine's MBean can be used to get a list of transaction identities (xid) and to commit and roll back transactions:

- `getPreparedTransactions()`
- `commitPreparedTransaction(String xid)`
- `rollbackPreparedTransaction(String xid)`

- b. To invoke the methods, you can use a `wsadmin` command, for example, you can use a command of the following form to obtain a list of the indoubt transaction identities from a messaging engine's MBean:

```
wsadmin> $AdminControl invoke [$AdminControl queryNames type=SIBMessagingEngine,*] getPreparedTransactions
```

Alternatively, you can use a script such as the following to invoke the methods on the MBean:

```
foreach mbean [$AdminControl queryNames type=SIBMessagingEngine,*] {
  set input 0

  while {$input >=0} {
    set xidList [$AdminControl invoke $mbean getPreparedTransactions]

    set meCfgId [$AdminControl getConfigId $mbean]
    set endIndex [expr {[string first "(" $meCfgId] - 1}]
    set me [string range ${meCfgId} 0 $endIndex]

    puts "----Prepared Transactions for ME $me ----"
    set index 0
    foreach xid $xidList {
      puts "  Index=$index XID=$xid"
      incr index
    }
    puts "----- End of list -----"
    puts "Select index of XID to commit/rollback (-1 to continue) : "
    set input [gets stdin]

    if {$input < 0} {
      puts "No index selected, going to continue."
    } else {
      set xid [lindex $xidList $input]
      puts "Enter c to commit or r to rollback XID $xid"
      set input [gets stdin]
      if {$input == "c"} {
        puts "Committing xid=$xid"
        $AdminControl invoke $mbean commitPreparedTransaction $xid
      }
      if {$input == "r"} {
        puts "Rolling back xid=$xid"
        $AdminControl invoke $mbean rollbackPreparedTransaction $xid
      }
    }
  }
}
```

```

    }
    puts ""
  }
}

```

This script lists the transaction identities of the transactions together with an index. You can then select an index and commit or roll back the transaction corresponding to that index.

Results

In summary, to identify and resolve indoubt transactions:

1. Use the administrative console to find the transaction identity of indoubt transactions.
2. If a transaction identity appears in the transaction management panel, commit or roll back the transactions as required.
3. If a transaction identity does not appear in the transaction management panel, use the methods on the messaging engine's MBean. For example, use a script to display a list of transaction identities for indoubt transactions. For each transaction:
 - a. Enter the index of the transaction identity of the transaction.
 - b. Enter c to commit the transaction
 - c. Enter r to roll back the transaction.
4. To check that transactions are no longer indoubt, restart the server and use the transaction management panel, or methods on the messaging engine's MBean.

Related tasks

“Viewing the service integration bus” on page 70

Use the Service Integration Bus browser on the administrative console to view the service integration bus.

Reviewing DB2 diagnostic information

Use a text editor to view the DB2[®] diagnostic log file on the machine where you suspect a problem to have occurred. The most recent events recorded are the furthest down the file.

About this task

Review DB2 diagnostic information when your systems are not working well. This is a way to see if the log files are full.

Procedure

On Unix type the following command: `tail -f /home/db2inst1/sqllib/db2dump/db2diag.log`

If the database is unresponsive, you will see something similar to the following:

```

2008-04-03-11.57.18.988249-300 I1247882009G504    LEVEL: Error
PID       : 16020                               TID  : 3086133792  PROC : db2agent (WPRCSDB) 0
INSTANCE: db2inst1                             NODE : 000        DB   : WPRCSDB
APPHDL   : 0-658                               APPID: 9.5.99.208.24960.080403084643
AUTHID   : DB2INST1
FUNCTION: DB2 UDB, data protection services, sqlpWriteLR, probe:6680
RETCODE  : ZRC=0x85100009=-2062548983=SQLP_NOSPACE
          "Log File has reached its saturation point"
          DIA8309C Log file was full.

```

```

2008-04-03-11.57.18.994572-300 E1247882514G540    LEVEL: Error

```

```
PID      : 16020                TID : 3086133792  PROC : db2agent (WPRCSDB) 0
INSTANCE: db2inst1           NODE : 000        DB   : WPRCSDB
APPHDL  : 0-658              APPID: 9.5.99.208.24960.080403084643
AUTHID   : DB2INST1
FUNCTION: DB2 UDB, data protection services, sqlpgResSpace, probe:2860
MESSAGE : ADM1823E The active log is full and is held by application handle
         "274". Terminate this application by COMMIT, ROLLBACK or FORCE
         APPLICATION.
```

In the preceding example, looking at the DB line, you can see that the WPRCSDB is experiencing full transaction logs.

Another way of viewing the db2diag logs is to log in as the DB2 user and run db2diag:

```
su -l db2inst1
db2diag | less
```

Related information

[↗ Interpreting diagnostic log file entries](#)

Process recovery troubleshooting tips

Using Business Process Choreographer Explorer can facilitate process recovery efforts.

The Business Process Choreographer Explorer provides a user interface for administrators to manage business processes and human tasks.

You can use the Business Process Choreographer Explorer to check the status of the Business Process Choreographer database (BPEDB). If you are unable to retrieve database information through the Business Process Choreographer Explorer, or if the Business Process Choreographer is slow to return database information, it might be an indication of a problem with the database.

Attempting to retrieve thousands of process instances or tasks is not wise if performance or database problems are suspected. Selecting a view which does not retrieve considerable data, such as “My Process Templates”, or limiting the amount of data retrieved for another view would be better options.

Related information

- [↗ Repairing processes and activities](#)
- [↗ Configuring Business Process Choreographer Explorer](#)
- [↗ Starting Business Process Choreographer Explorer](#)
- [↗ Business Process Choreographer Explorer overview](#)
- [↗ Tuning Business Process Choreographer Explorer](#)

About recovering the messaging subsystem

If the messaging system experiences problems you may need to recover the underlying messaging subsystem.

Typically this involves checking the state of various queues but can also include analyzing the integration bus infrastructure.

Detailed information on recovering the messaging subsystem can be found in the WebSphere Application Server information center.

Related concepts



Enterprise service bus messaging infrastructure

WebSphere Process Server includes enterprise service bus capabilities. WebSphere Process Server supports the integration of service-oriented, message-oriented, and event-driven technologies to provide a standards-based, messaging infrastructure in an integrated enterprise service bus.

Related information



Troubleshooting service integration message problems

IBM Support Assistant

The IBM Support Assistant is a tool that helps you use various IBM Support resources.

Note: The IBM Support Assistant is supported on Microsoft® Windows and Linux® systems.

The IBM Support Assistant offers four components to help you with software questions:

- a Search component, which helps you access pertinent Support information in multiple locations.
- a Product Information component, which helps you find the right IBM site for your product questions.
- a Tools component, which provides specialized analysis tools to investigate product problems.
- a Service component, which helps you submit an enhanced problem report that includes key system data to IBM.

Using the IBM Support Assistant with WebSphere Process Server, requires installing IBM Support Assistant and then installing plug-ins for WebSphere Process Server. The plug-ins for WebSphere Process Server include an automated way to gather information about a problem and send it to IBM, and tools that help you set trace levels.

For more information and to install the latest version of IBM Support Assistant, see the IBM Support Assistant Web page.

IBM Support Assistant also is included on the *WebSphere Application Server Network Deployment Supplements v7.0* disks that are included with WebSphere Process Server.

After the IBM Support Assistant is installed, you can start it with the **Start** menu option on Windows operating systems or with the `startisa.sh` shell script on all other platforms. On Windows operating systems, the IBM Support Assistant opens in its own window. On all other platforms, it opens in a Web browser.

When you have IBM Support Assistant open, you can view available plug-ins for WebSphere Process Server by clicking **Updater**, clicking **New Plug-ins** and then expanding **WebSphere**. When you select the check box for the WebSphere Process Server plug-in, and click **Install**, the download page opens.

To learn more about how to use the IBM Support Assistant, click **Help** in the IBM Support Assistant window.

Related tasks

“Getting fixes” on page 89

A product fix might be available to resolve your problem.

“Searching knowledge bases” on page 85

You can often find solutions to problems by searching IBM knowledge bases. Optimize your results by using available resources, support tools, and search methods.

“Contacting IBM Software Support” on page 91

IBM Software Support provides assistance with product defects.

Related reference

“Contacting IBM Software Support” on page 91

IBM Software Support provides assistance with product defects.

 [IBM Support Assistant](#)

Searching knowledge bases

You can often find solutions to problems by searching IBM knowledge bases. Optimize your results by using available resources, support tools, and search methods.

About this task

To search for solutions to your problems in IBM knowledge bases, perform the following steps.

Procedure

1. Search with IBM Support Assistant. IBM Support Assistant (ISA) is a free software serviceability workbench that helps you resolve questions and problems with IBM software products. The ISA tool can search multiple knowledge bases simultaneously

To search multiple Internet resources for your product, open the ISA and click **Search**. From this page, you can search a variety of resources including:

- IBM Software Support Documents
- IBM developerWorks®
- IBM newsgroups and forums
- Google
- IBM product information centers

Note: These free newsgroups and forums do not offer any formal IBM product support. They are intended for user-to-user communication. IBM will not be actively participating in these discussions. However, IBM does review these newsgroups periodically to maintain a free flow of accurate information. You may also want to browse the following resources individually.

2. Search with the IBM Software Support Toolbar. IBM Software Support Toolbar is a browser plug-in that provides you with a mechanism to easily search IBM support sites.
3. Search the information center.
IBM provides extensive documentation in the form of online information centers. An information center can be installed on your local machine or on a local intranet. An information center can also be viewed on the IBM Web site. You can use the powerful search function of the information center to query conceptual and reference information and detailed instructions for completing tasks.
4. Search available technical resources. In addition to this information center, the following technical resources are available to help you answer questions and resolve problems:
 - WebSphere Process Server technotes
 - WebSphere Process Server Authorized Program Analysis Reports (APARs)
 - WebSphere Process Server support Web site
 - WebSphere Redbooks® Domain
 - IBM Education Assistant
 - WebSphere Process Server forums and newsgroups








What to do next

Tip:

The following resources describe how to optimize your search results:

- Searching the IBM Support Web site
- Using the Google search engine
- IBM Software Support RSS feeds
- My Support e-mail updates

Related reference

-  [IBM WebSphere Process Server technotes](#)
 -  [IBM WebSphere Process Server Authorized Program Analysis Reports \(APARs\)](#)
 -  [IBM WebSphere Process Server support Web site](#)
 -  [IBM WebSphere Redbooks Domain](#)
 -  [IBM Education Assistant](#)
 -  [WebSphere Process Server forums and newsgroups](#)
- “IBM Support Assistant” on page 83
The IBM Support Assistant is a tool that helps you use various IBM Support resources.
-  [IBM Software Support Toolbar](#)

IBM Support Assistant

The IBM Support Assistant is a tool that helps you use various IBM Support resources.

Note: The IBM Support Assistant is supported on Microsoft Windows and Linux systems.

The IBM Support Assistant offers four components to help you with software questions:

- a Search component, which helps you access pertinent Support information in multiple locations.
- a Product Information component, which helps you find the right IBM site for your product questions.
- a Tools component, which provides specialized analysis tools to investigate product problems.
- a Service component, which helps you submit an enhanced problem report that includes key system data to IBM.

Using the IBM Support Assistant with WebSphere Process Server, requires installing IBM Support Assistant and then installing plug-ins for WebSphere Process Server. The plug-ins for WebSphere Process Server include an automated way to gather information about a problem and send it to IBM, and tools that help you set trace levels.

For more information and to install the latest version of IBM Support Assistant, see the IBM Support Assistant Web page.

IBM Support Assistant also is included on the *WebSphere Application Server Network Deployment Supplements v7.0* disks that are included with WebSphere Process Server.

After the IBM Support Assistant is installed, you can start it with the **Start** menu option on Windows operating systems or with the `startisa.sh` shell script on all other platforms. On Windows operating systems, the IBM Support Assistant opens in its own window. On all other platforms, it opens in a Web browser.

When you have IBM Support Assistant open, you can view available plug-ins for WebSphere Process Server by clicking **Updater**, clicking **New Plug-ins** and then expanding **WebSphere**. When you select the check box for the WebSphere Process Server plug-in, and click **Install**, the download page opens.

To learn more about how to use the IBM Support Assistant, click **Help** in the IBM Support Assistant window.

Related tasks

“Getting fixes” on page 89

A product fix might be available to resolve your problem.

“Searching knowledge bases” on page 85

You can often find solutions to problems by searching IBM knowledge bases. Optimize your results by using available resources, support tools, and search methods.

“Contacting IBM Software Support” on page 91

IBM Software Support provides assistance with product defects.

Related reference

“Contacting IBM Software Support” on page 91

IBM Software Support provides assistance with product defects.

 [IBM Support Assistant](#)

Getting fixes

A product fix might be available to resolve your problem.

About this task

To get product fixes, perform the following steps.

Procedure

1. Determine which fix you need. Check the list of WebSphere Process Server recommended fixes to confirm that your software is at the latest maintenance level. Check the list of problems fixed in the IBM WebSphere Process Server fix readme documentation that is available for each listed fix pack and refresh pack to see if IBM has already published an individual fix to resolve your problem. To determine what fixes are available using IBM Support Assistant, run a query on fix from the search page.

Individual fixes are published as often as necessary to resolve defects in WebSphere Process Server. In addition, two kinds of cumulative collections of fixes, called fix packs and refresh packs, are published periodically for WebSphere Process Server, in order to bring users up to the latest maintenance level. You should install these update packages as early as possible in order to prevent problems.

Note: Fixes specific to the underlying WebSphere Application Server product may also be obtained from the WebSphere Application Server Support Site or from the WebSphere Application Server Support team. Fixes for individual APARs for WebSphere Application Server generally can be applied without affecting WebSphere Process Server. However, consult with the software requirements page before updating WebSphere Application Server with cumulative collections of fixes (fix packs). First check to see that the cumulative fix has passed certification, or contact the Support team for verification.

2. Download the fix. Open the download document and follow the link in the **Download package** section. When downloading the file, ensure the name of the maintenance file is not changed. This includes both intentional changes and inadvertent changes caused by certain web browsers or download utilities.
3. Apply the fix. Follow the instructions in the **Installation Instructions** section of the download document. For more information, see the "Installing fix packs and refresh packs with the Update Installer" topic in the Installing WebSphere Process Server documentation.
4. Optional: To receive weekly notification of fixes and updates, subscribe to My Support e-mail updates.

Related reference

"IBM Support Assistant" on page 83

The IBM Support Assistant is a tool that helps you use various IBM Support resources.



Subscribe to My Support e-mail updates



Recommended Fixes for WebSphere Process Server

Contacting IBM Software Support

IBM Software Support provides assistance with product defects.

Before you begin

To take advantage of unique Support features, see the WebSphere Process Server support page. The Support Page contains the latest information on fixes and downloads, educational resources, and commonly encountered problems and their solutions.

Before contacting IBM Software Support, your company must have an active IBM software subscription and support contract, and you must be authorized to submit problems to IBM. The type of software subscription and support contract that you need depends on the type of product you have. For information about the types of software subscription and support contracts available, see "Enhanced Support" in the *Software Support Handbook* site listed in the Related Topics section.

To contact IBM Software Support with a problem, perform the following steps.

Procedure

1. Define the problem, gather background information, and determine the severity of the problem. For help, see the "Contacting IBM" in the *Software Support Handbook*.
2. Gather diagnostic information. When explaining a problem to IBM, be as specific as possible. Include all relevant background information so that IBM Software Support specialists can help you solve the problem efficiently. For information that IBM Support needs in order to help you solve a problem, see the WebSphere Process Server MustGather technote.

Tip: You can use the WebSphere Process Server plug-in for the IBM Support Assistant to capture data and send it to IBM.

Note: If you are able to determine that the problem is purely with underlying WebSphere Application Server functionality, consider requesting assistance specifically from the WebSphere Application Server Support team rather than the WebSphere Process Server team. For information that IBM Support needs in order to help you solve a WebSphere Application Server problem, see the WebSphere Application Server MustGather Technote.

3. Submit your problem to IBM Software Support in one of the following ways:
 - Using IBM Support Assistant: See the "IBM Support Assistant" topic.
 - Online: Open a service request on the IBM Software Support site using the Electronic Service Request (ESR) tool.
 - By telephone: For the telephone number to call in your country or region, go to the contacts page of the IBM Software Support Handbook on the Web and click the name of your geographic region.

What to do next

If the problem you submit is for a software defect or for missing or inaccurate documentation, IBM Software Support creates an Authorized Program Analysis

Report (APAR). The APAR describes the problem in detail. Whenever possible, IBM Software Support provides a workaround that you can implement until the APAR is resolved. Support will work and communicate with you on the progress and deliver the fix once it is completed. Additionally, once completed, IBM will also publish the resolved APARs on the Software Support Web site, so that other users who experience the same problem can benefit from the same resolution.

Related tasks

“IBM Support Assistant” on page 83

The IBM Support Assistant is a tool that helps you use various IBM Support resources.

Related reference

 [WebSphere Process Server Support](#)

 [Software Support Handbook](#)

 [MustGather: Read first for WebSphere Process Server for Version 6](#)

 [MustGather: Read first for all WebSphere Application Server products](#)

“IBM Support Assistant” on page 83

The IBM Support Assistant is a tool that helps you use various IBM Support resources.

 [IBM Software Support site](#)



Printed in USA