WebSphere® Process Server for Multiplatforms

Version 6.1.2

Troubleshooting and Support

WebSphere® Process Server for Multiplatforms

Version 6.1.2

Troubleshooting and Support

Note

Before using this information, be sure to read the general information in the Notices section at the end of this document.

27 June 2008

This edition applies to version 6, release 1, modification 2 of WebSphere Process Server for Multiplatforms (product number 5724-L01) 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.

Contents

Chapter 1. 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 nor 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?″ This 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, crash, 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 easy, 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 easily 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 easy 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 easier 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?

Chapter 2. 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 Chapter 14, “Contacting IBM Software Support,” on page 89.

Chapter 3. 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 tasks

Messages

This reference information provides additional information about messages you can encounter while using the product. It is organized according to the identifier of the product feature that produces the message.

Messages: installation and profile creation

Some of the most commonly found error messages encountered when installing and configuring can be addressed with actions that resolve the underlying problems.

Related reference

"Messages: installation and profile creation" on page 14
Some of the most commonly found error messages encountered when installing and configuring can be addressed with actions that resolve the underlying problems.

Chapter 4. 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 comprise 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 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. This set of logs is described in ″Profile-specific log files″ topic.

Related tasks


Installation and profile creation log files

Various log files are created during installation and uninstallation of WebSphere Process Server and during profile creation, augmentation, and deletion. Consult the applicable logs if problems occur during these procedures.

Related reference

“Profile-specific log files” on page 58
There are log files detailing the characteristics and runtime activities of individual profiles. These log files are located within the profile directory for each profile.


Installation and profile creation log files

Various log files are created during installation and uninstallation of WebSphere Process Server and during profile creation, augmentation, and deletion. Consult the applicable logs if problems occur during these procedures.

Chapter 5. Troubleshooting installation and configuration

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

About this task

The installer program records the following indicators of success at the end of the primary log file, which can be found in *install_root*/logs/wbi/install/log.txt on i5/OS®, Linux® and UNIX® or *install_root*\logs\wbi\install\log.txt on Windows®, where *install_root* represents the product installation directory:

- `INSTCONFSUCCESS`: installation was successful
- `INSTCONFPARTIALSUCCESS`: installation was partly successful. Some installation actions failed but can be retried.
- `INSTCONFFAILED`: installation was not successful. Recovery is not possible.

If the result is `INSTCONFPARTIALSUCCESS` or `INSTCONFFAILED`, continue analyzing the problem by following these steps.

To troubleshoot the installation, perform the following steps.

Procedure

1. Read any error messages from the installation process.

   See the following topic for an explanation: Error messages: installation and profile creation and augmentation. If the message corresponds to any of those described, correct the problem, clean up any installed portions, and try to reinstall.

   For details on uninstalling any installed portions before reinstalling, see Preparing for reinstallation after a failed uninstallation.
2. Determine if the installation of WebSphere Application Server Network Deployment was successful. (If it was unsuccessful and WebSphere Application Server Network Deployment was installed as part of the WebSphere Process Server installation, the installation process will not continue and an error message will be displayed.) If the installation of WebSphere Process Server was not successful, first check *install_root*/logs/install/log.txt on Linux and UNIX platforms or *install_root*\logs\install\log.txt on Windows platforms for errors to determine if the installation of WebSphere Application Network Deployment was successful.

   If the installation of WebSphere Application Server Network Deployment failed, see the topic Troubleshooting installation in the WebSphere Application Server Network Deployment information center and use the information found there to correct the problems before attempting to reinstall WebSphere Process Server.

   If the installation of WebSphere Application Server Network Deployment succeeded and the installation of WebSphere Process Server failed, use the troubleshooting information below to correct the problems.

   i5/OS On i5/OS platforms, if the installation was unsuccessful and WebSphere Application Server Network Deployment was installed as part of the WebSphere Process Server installation, the installation process will not continue and an error message will be displayed. If the installation of WebSphere Process Server was not successful, first check *install_root*/logs/

install/log.txt on i5/OS platforms for errors to determine if the installation of WebSphere Application Network Deployment was successful.

3. Check the WebSphere Process Server installation log files for errors after installing.

   For information about the names, locations, and descriptions of the various log files that are created, see Installation and profile creation log files.

   Check the log files in this sequence:

   i5/OS **On i5/OS platforms:**

   a. log files in the *install_root*/logs/wbi/install directory
   b. *user_data_root*/profileRegistry/logs/manageprofiles/*profile_name*_create.log, *user_data_root*/profileRegistry/logs/manageprofiles/*profile_name*_augment.log, and *user_data_root*/profileRegistry/logs/manageprofiles/pmt.log.
   c. Any additional log or trace files generated by installation actions. Look in *install_root*/logs/wbi/install for trace files generated during the installation process. Look in *install_root* /logs/manageprofiles/*profile_name* for those generated by profile creation or augmentation. (For more information about *install_root* and *profile_root* locations, see Default installation directories for the product, profiles, and tools.) These files are primarily intended for use by IBM technical support.

   Linux UNIX **On Linux and UNIX platforms:**

   a. log files in the *install_root*/logs/wbi/install directory
   b. log files in the %tmp%/niflogs.wbi directory if no files are found in *install_root*/logs/wbi/install
   c. *install_root*/logs/manageprofiles/*profile_name*_create.log, *install_root*/logs/manageprofiles/*profile_name*_augment.log, and *install_root*/logs/manageprofiles/pmt.log.
   d. Any additional log or trace files generated by installation actions. Look in *install_root*/logs/wbi/install for trace files generated during the installation process. Look in *install_root*/logs/manageprofiles/*profile_name* for those generated by profile creation or augmentation. (For more information about *install_root* and *profile_root* locations, see Default installation directories for the product, profiles, and tools.) These files are primarily intended for use by IBM technical support.

   Windows **On Windows platforms:**

   a. log files in the *install_root*\logs\wbi\install directory
   b. log files in the %tmp%\niflogs.wbi directory if no files are found in *install_root*\logs\wbi\install
   c. *install_root*\logs\manageprofiles*profile_name*_create.log, *install_root*\logs\manageprofiles*profile_name*_augment.log, and *install_root*\logs\manageprofiles\pmt.log.
   d. Any additional log or trace files generated by installation actions. Look in *install_root*\logs\wbi\install for trace files generated during the installation process. Look in *install_root*\logs\manageprofiles*profile_name* for those generated by profile creation or augmentation. (For more information about *install_root* and *profile_root* locations, see Default installation directories for the product, profiles, and tools.) These files are primarily intended for use by IBM technical support.

4. Determine whether the installation problem is caused by a configuration script that failed.

The *install_root*/logs/wbi/installconfig.log file on i5/OS, Linux and UNIX platforms or *install_root*\logs\wbi\installconfig.log file on Windows platforms indicates configuration problems that can prevent the product from working correctly.

For more information about diagnosing failed configuration scripts, see "Diagnosing a failing Ant configuration script" on page 23.

5. If the error logs do not contain enough information to determine the cause of the problem, uninstall the product, clean up any log files or other artifacts that are left behind, turn on tracing, and reinstall.
   - Report the stdout and stderr logs to the console window by adding the **-is:javaconsole** parameter to the install command:
     - i5/OS **On i5/OS platforms:**

       ```
       install -is:javaconsole
       ```

       Capture the stream to a file with the following commands:

       ```
       install -is:javaconsole > captureFileName.txt 2>&1
       ```

     - Linux UNIX **On Linux and UNIX platforms:**

       ```
       install -is:javaconsole
       ```

       Capture the stream to a file with the following commands:

       ```
       install -is:javaconsole > captureFileName.txt 2>&1
       ```

     - Windows **On Windows platforms:**

       ```
       install.bat -is:javaconsole
       ```

       Capture the stream to a file with the following commands:

       ```
       install.bat -is:javaconsole > drive:\captureFileName.txt
       ```

   - Capture additional information to a log of your choice with the `-is:log file_name` option.
6. If you have successfully created a server profile, use the First steps console or the command-line method to start the server. For more information, see First steps console.
7. Verify that the server starts and loads properly by looking for a running Java™ process and the *Open for e-business* message in the SystemOut.log and SystemErr.log files.

   If no Java process exists or if the message does not appear, examine the same logs for any miscellaneous errors. Correct any errors and retry.

   You can find the SystemOut.log and SystemErr.log files in the following platform-specific directories:

   - i5/OS **On i5/OS platforms:** *profile_root*/logs/*servername*
   - Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs/*servername*
   - Windows **On Windows platforms:** *profile_root*\logs*servername*
8. Use the First steps console or the command-line method to stop the server, if it is running. For more information, see Options on the First steps console
9. To troubleshoot a WebSphere Process Server deployment environment, see Verifying your deployment environment.
10. If you want to use a Snoop Servlet to verify the ability of the Web server to retrieve an application from WebSphere Process Server, see step 10 in Troubleshooting installation in the WebSphere Application Server Network Deployment documentation.
11. Start the administrative console. For more information, see Starting and stopping the administrative console.

12. To resolve any IP address caching problems, see step 14 in Troubleshooting installation in the WebSphere Application Server Network Deployment documentation.

What to do next

On the product support Web site, you can review current information about resolutions to known problems, and you can read documents that can save you time gathering the information that you need to resolve a problem. Before opening a PMR, see the IBM WebSphere Process Server support page.

Messages: installation and profile creation

Some of the most commonly found error messages encountered when installing and configuring can be addressed with actions that resolve the underlying problems.

Note: Linux UNIX Windows The following WebSphere Process Server installation and configuration errors appear on Linux, UNIX, and Windows platforms.

Tip: For information about messages that might be generated by the installation of WebSphere Application Server Network Deployment, refer to the Messages issued during installation and profile creation topic available in the WebSphere Application Server Network Deployment documentation.

What kind of problem are you having while installing WebSphere Process Server?

- "Supported IBM JDK was not found. The IBM JDK shipped with this product must be located at *install_root*/JDK. Please correct this problem and try again." on page 15
- "Warning: Cannot convert string "<type_name>"to type FontStruct" on page 15

If you do not see an error message that resembles yours, or if the information provided does not solve your problem, contact WebSphere Process Server support at IBM for further assistance.

Related tasks

Messages

This reference information provides additional information about messages you can encounter while using the product. It is organized according to the identifier of the product feature that produces the message.

Chapter 3, "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.

Related reference

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.

Supported IBM JDK was not found. The IBM JDK shipped with this product must be located at *install_root*/JDK. Please correct this problem and try again.

If you use symbolic links to point to the IBM Java Development Kit (JDK) shipped with the product, or to a JDK found in the PATH environment variable on your system, IBM SDK for Java validation might fail, resulting in a failed installation. This problem is caused by the way IBM SDK for Java validation code detects whether the JDK shipped with the product is the current JDK used for installation.

To resolve this problem, do not use symbolic links in JVMs supplied with the installation image of WebSphere Process Server and remove symbolic links from all JVMs that appear in your system's PATH environment variable.

Warning: Cannot convert string "<type_name>"to type FontStruct

If you install the Web server plug-ins for WebSphere Application Server, you also install the ikeyman utility. The ikeyman utility is part of the Global Services Kit 7 (GSKit7).

Linux If you issue the ikeyman.sh script on a Linux system, you might see the following message:

```
Warning: Cannot convert string
   "-monotype-arial-regular-r-normal--*-140-*-*-p-*-iso8859-1"
   to type FontStruct
```

You can safely ignore the warning and use the ikeyman utility.

Installation and profile creation log files

Various log files are created during installation and uninstallation of WebSphere Process Server and during profile creation, augmentation, and deletion. Consult the applicable logs if problems occur during these procedures.

Table 1 on page 16 shows the logs, content, and indicators of success and failure for WebSphere Process Server.

If the logs directory does not exist on your system, the installation failed very early in the process. In this case, review the following:

- Linux UNIX /tmp/niflogs.wbi/log.txt file on Linux and UNIX platforms.
- Windows %TEMP%\niflogs.wbi\log.txt file on Windows platforms.
- i5/OS /tmp/InstallShield/niflogs.wbi/log.txt file on i5/OS platforms.

Important: Windows **On Windows platforms**: The %TEMP% directory can be hidden from the Windows GUI. It usually resolves to C:\Documents and Settings*username*\Local Settings\Temp. To find the %TEMP% directory, type one of the following commands at a command prompt:

- At a command prompt, type `cd %TEMP%`.
- At a command prompt, type `echo %TEMP%` and copy and paste the result into Windows Explorer.

Some directory paths, file names, and indicator values in Table 1 contain spaces to allow the entries to fit in the table cells. The actual directory paths, file names, and indicator values do not contain spaces.

The variable *install_root* represents the installation directory of WebSphere Process Server. The variable *profile_root* represents the root location of a profile.

i5/OS **On i5/OS platforms:** The variable *user_data_root* represents the default user data directory.

For more information see Default installation directories for the product, profiles, and tools.

Table 1. Installation and profile logs for WebSphere Process Server components

Log	Content	Indicators
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/wbi/install/log.txt • Windows **On Windows platforms:** *install_root*\logs\wbi\install\log.txt • i5/OS **On i5/OS platforms:** *install_root*/logs/wbi/install/log.txt	Logs all installation events relating to WebSphere Process Server.	**INSTCONFFAILED** Total installation failure. **INSTCONFSUCCESS** Successful installation. **INSTCONFPARTIALSUCCESS** Installation errors occurred but the installation is still usable. Additional information in other log files identifies the errors.
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/wbi/installconfig.log • Windows **On Windows platforms:** *install_root*\logs\wbi\installconfig.log • i5/OS **On i5/OS platforms:** *install_root*/logs/wbi/installconfig.log	Logs configuration actions that run at the end of the installation process to configure components, install system applications, and create Windows shortcuts and registry entries.	Contains a series of <record> elements that document the configuration actions. If a post-installation configuration action fails, text like the following appears in the log: <record> <date>2005-05-26T11:41:17</date> <millis>1117132877344</millis> <sequence>742</sequence> <logger>com.ibm.ws.install.configmanager.ConfigManager</logger> <level>WARNING</level> <class>com.ibm.ws.install.configmanager.ConfigManager</class> <method>executeAllActionsFound</method> <thread>12</thread> <message>Configuration action failed: com.ibm.ws.install.configmanager.actionengine.ANTAction-D:\WBI\AS\properties\version\install.wbi\6.1.0.0\config\full\install\90SInstallCEI.ant</message> </record> If no actions fail, the following message is included in the record in the log: <record> . . . <message>No errors were encountered while executing the repository actions</message> </record>

Table 1. Installation and profile logs for WebSphere Process Server components (continued)

Log	Content	Indicators
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/manageprofiles/pmt.log • Windows **On Windows platforms:** *install_root*\logs\manageprofiles\pmt.log • i5/OS **On i5/OS platforms:** *user_data_root*/profileRegistry/logs/manageprofiles/pmt.log	Logs all events from the Profile Management Tool.	**INSTCONFFAILED** Total profile creation failure. **INSTCONFSUCCESS** Successful profile creation. **INSTCONFPARTIALSUCCESS** Profile creation errors occurred but the profile is still functional. Additional information in other log files identifies the errors.
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/manageprofiles/*profile_name*_create.log • Windows **On Windows platforms:** *install_root*\logs\manageprofiles*profile_name*_create.log • i5/OS **On i5/OS platforms:** *user_data_root*/profileRegistry/logs/manageprofiles/*profile_name*_create.log	• Traces all events that occur during the creation of the named profile. • Created when a profile is created during a Complete installation, when using the Profile Management Tool, or when using the manageprofiles command.	**INSTCONFFAILED** Total profile creation failure. **INSTCONFSUCCESS** Successful profile creation. **INSTCONFPARTIALSUCCESS** Profile creation errors occurred but the profile is still functional. Additional information in other log files identifies the errors.
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/manageprofiles/*profile_name*_augment.log • Windows **On Windows platforms:** *install_root*\logs\manageprofiles*profile_name*_augment.log • i5/OS **On i5/OS platforms:** *user_data_root*/profileRegistry/logs/manageprofiles/*profile_name*_augment.log	• Traces all events that occur during the augmentation of the named profile. • Created when a profile is augmented, when using the Profile Management Tool, or when using the manageprofiles command.	**INSTCONFFAILED** Total profile augmentation failure. **INSTCONFSUCCESS** Successful profile augmentation. **INSTCONFPARTIALSUCCESS** Profile augmentation errors occurred but the profile is still functional. Additional information in other log files identifies the errors.

Table 1. Installation and profile logs for WebSphere Process Server components (continued)

Log	Content	Indicators
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/manageprofiles/*profile_name*_delete.log • Windows **On Windows platforms:** *install_root*/logs/manageprofiles/*profile_name*_delete.log • i5/OS **On i5/OS platforms:** *user_data_root*/profileRegistry/logs/manageprofiles/*profile_name*_delete.log	• Traces all events that occur during the deletion of the named profile. • Created when profile deletion is performed with the manageprofiles command.	**INSTCONFFAILED** Total profile deletion failure. **INSTCONFSUCCESS** Successful profile deletion. **INSTCONFPARTIALSUCCESS** Profile deletion errors occurred but the profile is still deleted. Additional information in other log files identifies the errors.
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/install/log.txt • Windows **On Windows platforms:** *install_root*\logs\install\log.txt • i5/OS **On i5/OS platforms:** *install_root*/logs/wbi/install/log.txt	• Logs all installation events relating to WebSphere Application Server Network Deployment. • Created as part the underlying installation of WebSphere Application Server Network Deployment that is installed with WebSphere Process Server.	**INSTCONFFAILED** Total installation failure. **INSTCONFSUCCESS** Successful installation. **INSTCONFPARTIALSUCCESS** Installation errors occurred but the installation is still usable. Additional information in other log files identifies the errors.
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/installconfig.log • Windows **On Windows platforms:** *install_root*\logs\installconfig.log • i5/OS **On i5/OS platforms:** *install_root*/logs/wbi/installconfig.log	• Logs configuration actions that run at the end of the installation process to configure components, install system applications, and create Windows shortcuts and registry entries. • Created as part the underlying installation of WebSphere Application Server Network Deployment that is installed with WebSphere Process Server.	Contains a series of `<record>` elements that document the configuration actions.

Table 1. Installation and profile logs for WebSphere Process Server components (continued)

Log	Content	Indicators
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/wbi/uninstall/log.txt • Windows **On Windows platforms:** *install_root*\logs\wbi\uninstall\log.txt • i5/OS **On i5/OS platforms:** *install_root*/logs/wbi/uninstall/log.txt	Logs all uninstallation events relating to WebSphere Process Server.	**INSTCONFFAILED** Total uninstallation failure. **INSTCONFSUCCESS** Successful uninstallation. **INSTCONFPARTIALSUCCESS** The uninstallation wizard successfully removed the core product files, but errors occurred during configuration. Additional information in other log files identifies the errors.
• Linux UNIX **On Linux and UNIX platforms:** *install_root*/logs/wbi/update/updateconfig.log • Windows **On Windows platforms:** *install_root*\logs\wbi\update\updateconfig.log • i5/OS **On i5/OS platforms:** *install_root*/logs/wbi/update/updateconfig.log	Logs configuration actions that run at the end of the uninstallation process.	Contains a series of `<record>` elements that document the configuration actions.
• i5/OS **On i5/OS platforms:** %*TEMP*%\firststeps_i5.log	Logs errors that occur during the first steps procedure and provides suggestions on how to fix them.	This log is useful when you are running first steps from the command line, because there are chances of typos or similar errors. When first steps is launched from the Profile Management Tool or the installer, it should work without any problem. Whenever you experience any unexpected or erroneous behavior from first steps, you should check this log file.

Related tasks

Chapter 4, "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 comprise the second group.

Related reference

Chapter 4, "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 comprise the second group.

"Profile-specific log files" on page 58
There are log files detailing the characteristics and runtime activities of individual profiles. These log files are located within the profile directory for each profile.

Troubleshooting the launchpad application

If the launchpad application does not start, try the following troubleshooting tips.

Restart the launchpad after you make any changes.

- If you are using images from Passport Advantage®, make sure that you extract the contents of the images for *WebSphere Process Server V6.1.2 DVD*, *WebSphere Application Server Network Deployment Supplements V6.1* CD, and *WebSphere Application Server Toolkit V6.1.1 Disk 1* (if included for your platform) into three separate directories. Extracting the files from the images into the same directory will cause errors to occur. It is recommended that you use three sibling directories. For example, use a set of directories such as the following:

  Note: i5/OS The installation images obtained from Passport Advantage must be downloaded to a Windows workstation.

  – i5/OS

  ```
  %/downloads/WPS/image1
  %/downloads/WPS/image2
  %/downloads/WPS/image3
  ```

  – Linux UNIX

  ```
  %/downloads/WPS/image1
  %/downloads/WPS/image2
  %/downloads/WPS/image3
  ```

  – Windows

  ```
  C:\downloads\WPS\image1
  C:\downloads\WPS\image2
  C:\downloads\WPS\image3
  ```

- If you can start the launchpad, but selecting a link does not resolve to a page in the launchpad, you might have the media for the wrong operating system in the disk drive. Check the validity of the media.

- Windows If you are attempting to use the Mozilla browser on a Windows system, Internet Explorer might open instead. The launchpad does not recognize Mozilla as the default browser if Internet Explorer is also installed on the same system. The launchpad is fully functional with Internet Explorer, so no action is required.

  To create an environment variable that forces the use of Mozilla, issue the following case-specific command at a command prompt:

  ```
  set BROWSER=Mozilla
  ```

- Ensure that the JavaScript™ function is enabled in your browser.

  Linux UNIX Mozilla: Click **Edit > Preferences > Advanced > Scripts & Plugins**:

  – Enable JavaScript for: Navigator.

  – Allow scripts to ... (Select all boxes.)

  Linux UNIX Mozilla Firefox: Click **Tools > Options > Content**:

  – Select **Enable Java**.

  – Select **Enable JavaScript**.

  – Click **Advanced** and Allow scripts to ... (Select all boxes.)

  Windows Internet Explorer: Click **Tools > Internet Options > Security > Custom Level for Internet > Scripting > Active scripting > Enable**.

If the launchpad links still do not work after trying these tips, start the component installation programs directly. The locations of these programs are listed in Options on the launchpad.

Troubleshooting a silent installation

If a silent installation using a response file fails, you can examine log files and error messages to determine what went wrong, and make changes to your response file.

Before you begin

For information about using the response file for a silent installation of WebSphere Process Server, see Installing silentlyInstalling silently.

To troubleshoot a silent product installation, perform the following steps.

Procedure

1. Check your response file to make sure you are precise when supplying option values in the file so that the installation program can read the values. Incorrect specifications affect the silent interface of the installation wizard. For example, always use the correct case within property names, which are case-sensitive. In addition, always enclose values in double quotation marks. If the error is an incorrect option value, the InstallShield MultiPlatform program displays a warning message that you must confirm and stops the installation.
2. Compare your response file to the responsefile.wbis.txt file that is shipped with the product to make the necessary corrections. This file is in the *install_image*/WBI directory. After correcting the file, reinstall.
3. Review commonly found error messages in Messages: installation and profile creation and augmentation.
4. Examine log files. See the descriptions of relevant log files listed in Installation and profile creation log files.
5. Certain events can prevent InstallShield MultiPlatform from starting the installation wizard silently (for example, not having enough disk space to launch the installation wizard). If your installation fails and there is no information in the installation logs, record entries for events that cause the ISMP program to fail to start the installation wizard.

   The syntax of the install command for logging such events is as follows:

   AIX **On AIX® platforms:**

   ```
   install -options "/usr/IBM/WebSphere/silentFiles/myresponsefile.txt"
           -silent -log
   ```

   HP-UX Solaris **On HP-UX and Solaris platforms:**

   ```
   install -options "/opt/IBM/WebSphere/silentFiles/myresponsefile.txt"
           -silent -log
   ```

   i5/OS **On i5/OS platforms:**

   ```
   install -options responsefile.wbis.txt -silent -log log.txt @ALL
   ```

   Note: i5/OS **On i5/OS platforms:** You must change to the directory that contains the copied DVD image. Example: */MYDIR*/WBI

   Linux **On Linux platforms:**

   ```
   install -options "/opt/ibm/WebSphere/silentFiles/myresponsefile.txt"
           -silent -log
   ```

Windows **On Windows platforms:**

```
install.bat -options "C:\IBM\WebSphere\silentFiles\myresponsefile.txt"
            -silent -log # !C:\IBM\WebSphere\silentFiles\log.txt  @ALL
```

6. For other tips on troubleshooting your installation, see Troubleshooting installation.
7. If your profile did not create successfully, see Recovering from profile creation or augmentation failure.

i5/OS installation troubleshooting tips

You can refer to sources that might be helpful in troubleshooting an installation problem for a WebSphere Process Server product on the i5/OS operating system.

WebSphere Process Server offers several methods you can use to troubleshoot problems. Which method you use depends on the nature of the problem. Generally, you use a combination of these methods to determine the cause of a problem and then decide on an appropriate method for its resolution.

Tip 1: Refer to troubleshooting documentation for WebSphere Application Server for i5/OS

These resources provide general troubleshooting assistance:

- WebSphere Process Server Release Notes.
- WebSphere Application Server FAQ database.
- WebSphere Application Server for OS/400® newsgroup. This System i™ Technical Support Web-based forum is dedicated to WebSphere Application Server for i5/OS and OS/400.

Tip 2: Install WebSphere Process Server Version 6.1 for i5/OS

- **Wrong version of i5/OS installed on your server.**

  WebSphere Enterprise Service Bus runs on i5/OS V5R3 or V5R4. The product cannot be installed on prior releases of i5/OS. .
- **IBM Development Kit for Java V1.5 is not installed.**

  Local and remote command-line installations require JDK 1.5. Install product 5722-JV1, option 7 to obtain JDK 1.5. After installing option 7, you should reinstall the cumulative PTF package and Java group PTF to pick up any JDK 1.5 specific fixes.
- **Host servers are not started, or failed to start correctly.**

  The installation process requires that the i5/OS host servers be running. To start the host servers, run this command from the CL command line.

  ```
  STRHOSTSVR SERVER(*ALL)
  ```

  If errors other than "Host server daemon jobs unable to communicate using IPX." occur when starting the host servers, follow the instructions in the error message to fix the problem. After the problem is fixed, start the host servers and attempt to install WebSphere Process Server Server again.
- **Installation fails due to "Object not found" or "Not authorized" errors.**

  The user profile of the user installing the product must have *ALLOBJ and *SECADM special authorities.

Tip 3: Start WebSphere Process Server for i5/OS

- **Port conflicts**

Port conflicts may exist if you have multiple stand alone installs of WebSphere Application Server or multiple installs of stack products that bundle WebSphere Application Server like WebSphere Enterprise Service Bus or WebSphere Process Server on the same physical i5/OS machine.

Diagnosing a failing Ant configuration script

Determine whether a product installation problem on an operating system such as AIX, Linux, Windows, or i5/OS is caused by a failing Apache™ Ant configuration script.

Before you begin

Start diagnosing installation problems by looking at the troubleshooting procedure. See Troubleshooting installation. After the installation completes successfully, several Ant scripts configure the product. The following procedure describes what to do when an Ant script fails. When the installation log does not indicate a failure, determine how to correct any problems with failing Ant configuration scripts.

About this task

The *install_root*/logs/wbi/installconfig.log file, when present, describes any failure of an Ant script. Determine if any of the following configuration scripts failed. If so, use the configuration script recovery procedures. Use the investigative action to manually verify that the following configuration scripts ran successfully during the configuration of the WebSphere Process Server product. If any script failed, use the recovery action steps to complete the function of the script.

To diagnose failed Ant configuration scripts, perform the following steps.

- Diagnose the failed 90SConfigWBIMigrationScript.ant configuration script. This script changes the permissions of the following script to 755: *install_root*/bin/wbi_migration. This script also replaces the following tokens in the *install_root*/bin/wbi_migration script:

From:	To the value that you selected during installation:
`${JAVAROOT}`	*install_root*/java/jre/bin/java
`${MIGRATIONJAR}`	*install_root*/bin/migration/migrationGUI/migrationGUI.jar
`${WASROOT}`	*install_root*
`${PRODUCTID}`	`${WS_CMT_PRODUCT_TYPE}`

1. Investigative action: Verify that the permissions are 755 for the *install_root*/bin/wbi_migration.sh script on Linux and UNIX platforms, *install_root*\bin\wbi_migration.bat on Windows platforms, or for the *install_root*/bin/wbi_migration script on i5/OS platforms.
2. Recovery action: Issue the following command: `chmod 755 install_root/bin/wbi_migration.sh` on Linux and UNIX platforms, `chmod 755 install_root\bin\wbi_migration.bat` on Windows platforms or `chmod 755 install_root/bin/wbi_migration` on i5/OS platforms.
3. Investigative action: Open the *install_root*/bin/wbi_migration.sh on Linux and UNIX platforms, *install_root*\bin\wbi_migration.bat on Windows platforms, or *install_root*/bin/wbi_migration script on i5/OS platforms in an

editor and verify that real values exist instead of the following values: ${JAVAROOT}, ${MIGRATIONJAR}, ${WASROOT}, and ${PRODUCTID}.

4. Recovery action: Change the following tokens to actual values in the wbi_migration script: ${JAVAROOT}, ${MIGRATIONJAR}, ${WASROOT}, and ${PRODUCTID}.

- Diagnose the failed 85SConfigNoProfileFirstStepsWBI.ant. This script copies all files from the *install_root*/properties/version/install.wbi/firststeps.wbi directory to the *install_root*/firststeps/wbi/html/noprofile directory. This script also replaces the following tokens in the *install_root*/firststeps/wbi/firststeps.sh script (Linux, and UNIX), the *install_root*\firststeps\wbi\firststeps.bat script (Windows platforms), or the *install_root* /firststeps/wbi/firststeps script (i5/OS platforms):

From:	To the value that you selected during installation:
${JAVAROOT}	*install_root*/java/jre/bin/java
${PROFILEROOT}	*install_root*
${HTMLSHELLJAR}	*install_root*/lib/htmlshellwbi.jar
${CELLNAME}	${WS_CMT_CELL_NAME}

1. Investigative action: Verify that all files are copied from the *install_root*/properties/version/install.wbi/firststeps.wbi directory to the *install_root*/firststeps/wbi/html/noprofile directory.
2. Recovery action: Copy all of the files from the *install_root*/properties/version/install.wbi/firststeps.wbi directory to the *install_root*/firststeps/wbi/html/noprofile directory.
3. Investigative action: Open the *install_root*/firststeps/wbi/firststeps script in an editor. Verify that real values exist instead of the following values: ${JAVAROOT}, ${PROFILEROOT}, ${HTMLSHELLJAR}, and ${CELLNAME}.
4. Recovery action: Change the following tokens to actual values in the *install_root*/firststeps/wbi/firststeps script. ${JAVAROOT}, ${PROFILEROOT}, ${HTMLSHELLJAR}, and ${CELLNAME}.

Results

After you correct any installation errors and any Ant script configuration errors by performing the corrective actions in this procedure, the installation is complete.

What to do next

Start the First steps console.

Recovering from profile creation or augmentation failure

The Profile Management Tool can experience failures when creating new or augmenting existing profiles. The same can occur using the manageprofiles command. If such a failure occurs, first check the log files as described in this topic, then follow the recovery instructions described, depending on the situation.

Log files

All manageprofiles log files are in *install_root*/logs/manageprofiles. Look at the following log files in the order given. Each log file must contain the entry

"INSTCONFSUCCESS." If a file does not include this entry, a failure was detected. Look at the log files to determine why a failure was encountered and to determine a remedy.

1. The log file *profile_name*_create.log (where *profile_name* is the name of the profile).

   Note: Look at this file only if you were creating a new profile, not augmenting an existing one.

   This log file is located in the following directories:

   - i5/OS *user_data_root*/profileregistry/logs/manageprofiles on i5/OS systems (where *user_data_root* is the WebSphere Process Server user data directory).
   - Linux UNIX *install_root*/logs/manageprofiles on Linux and UNIX systems (where *install_root* is the WebSphere Process Server installation root location).
   - Windows *install_root*\logs\manageprofiles on Windows systems (where *install_root* is the WebSphere Process Server installation root location).

   Search for the text `Configuration action succeeded` or `Configuration action failed`.

   Note: There can be multiple occurrences of `Configuration action failed`. Investigate and remedy each one. Also review the log files described in the following options, if the profile was created.

   Note: Additional information is available in the manageprofiles directory in the pmt.log, which logs all events that occur when a default profile is created during a complete installation using the Profile Management Tool.
2. Log file *profile_name*_augment.log (where *profile_name* is the name of the profile).

   This log file is located in the following directories:

   - i5/OS *user_data_root*/profileregistry/logs/manageprofiles on i5/OS systems (where *user_data_root* is the WebSphere Process Server user data directory).
   - Linux UNIX *install_root*/logs/manageprofiles on Linux and UNIX systems (where *install_root* is the WebSphere Process Server installation root location).
   - Windows *install_root*\logs\manageprofiles on Windows systems (where *install_root* is the WebSphere Process Server installation root location).

   Search for the text `Configuration action succeeded` or `Configuration action failed`.

   Note: There can be multiple occurrences of `Configuration action failed`. Investigate and remedy each one. Also review the log files described in the following options, if the profile was created.
3. Individual profile template action log files.

   If you discovered false values in the log files described in the preceding options, review the log files in the following directories:

   - i5/OS *user_data_root*/profileregistry/logs on i5/OS systems.
   - Linux UNIX *install_root*/logs/manageprofiles/*profile_name* on Linux and UNIX systems.

- Windows *install_root*\logs\manageprofiles*profile_name* on Windows systems.

where *profile_root* or *user_data_root* is the installation location of the profile.

These log files do not follow a consistent naming convention, but typically, each is the name of the Ant script that failed followed by `.log`. For example, suppose the following entry is in the *profile_name*_augment.log file:

```
<messages>Result of executing
E:\o0536.15\profileTemplates\default.wbicore\actions\saveParamsWbiCore.ant
was:false</messages>
```

First look at the surrounding entries in the *profile_name*_augment.log file in the *install_root*/logs/manageprofiles directory. If you cannot determine the cause of the failure from the surrounding entries, look for the corresponding log file for any failing Ant script entries. In this case, the log file created by the `saveParamsWbiCore.ant` script is `saveParamsWbiCore.ant.log`. Look at that file to investigate why the failure occurred.

Recovery for creation failure

After you determine why profile creation failed and address the cause of the failure, you can try to create the profile again.

Note: When you create a profile, it first creates a WebSphere Application Server profile and then augments it with WebSphere Process Server profile templates to create a WebSphere Process Server profile. Even if you encountered a profile creation failure, a profile can exist that does not have all the needed augmentations.

To determine if the profile exists, run the *install_root*`/bin/manageprofiles -listProfiles` command. If the profile name you used for creation does not exist, you can recreate the profile. If the profile name you used for creation exists, then the profile was created and you have encountered an augmentation failure. For tips on recovering from an augmentation failure, see "Recovery for augmentation failure."

Recovery for augmentation failure

After you determine why profile augmentation failed and address the cause of the failure, you can try to augment the existing profile again to successfully create a complete WebSphere Process Server profile by following these steps:

1. Start the Profile Management Tool and, instead of creating a new profile, choose to augment an existing profile.
2. Choose the profile you were working with, and enter the correct information for it.

Note: Some of the augmentations might have completed successfully the first time you ran the Profile Management Tool. As a result, you might not see all of the panels that you saw the first time you tried to create the profile. This is because the Profile Management Tool detects which remaining augmentations must be completed and displays only the necessary panels.

Troubleshooting the Business Process Choreographer configuration

For information on how to solve problems relating to the configuration of Business Process Choreographer and its Business Flow Manager, or Human Task Manager components, go to the WebSphere Process Server for Multiplatforms, version 6.1, information center and review the topics under **Installing and configuring WebSphere Process Server > Troubleshooting installation and configuration > Troubleshooting the Business Process Choreographer configuration**. You can also find this information in the *Business Process Choreographer PDF*.

Chapter 6. Troubleshooting migration

If you encounter problems during migration, the information described here could help.

Troubleshooting version-to-version migration

Review this page for troubleshooting tips if you encounter problems while migrating from an older version of WebSphere Process Server.

- While you are using the version 6.1 migration wizard to create a profile before migrating a configuration, you might see the following profile-creation error messages.

  ```
  profileName: profileName cannot be empty
  profilePath: Insufficient disk space
  ```

  These error messages might be displayed if you enter a profile name that contains an incorrect character such as a space. Rerun the migration wizard, and verify that there are no incorrect characters in the profile name such as a space, quotes, or any other special characters.
- If you encounter a problem when you are migrating from a previous version of WebSphere Process Server to version 6.1, check your log files and other available information.
  1. Look for the log files, and browse them for clues.
     - *migration_backup_directory*/WBIPreUpgrade.*time_stamp*.log
     - *profile_root*/log/WASPostUpgrade.*time_stamp*.log
     - *install_root*/logs/clientupgrade.*time_stamp*.log
     - *profile_root*/logs/bpeupgrade.log
     - *migration_backup_directory*/WBIProfileUpgrade.ant.*timestamp*.log
  2. Look for `MIGR0259I: The migration has successfully completed.` or `MIGR0271W: The migration completed with warnings.` in the following directories:
     - *migration_backup_directory*/WBIPreUpgrade.*time_stamp*.log
     - *profile_root*/logs/WASPostUpgrade.*time_stamp*.log
     - *install_root*/logs/clientupgrade.*time_stamp*.log

     If `MIGR0286E: The migration failed to complete.` is displayed, attempt to correct any problems based on the error messages that appear in the log file. After correcting any errors, rerun the command from the bin directory of the product installation root.
  3. Open the Log and Trace Analyzer built into the Application Server Toolkit (AST) on the service log of the server that is hosting the resource that you are trying to access, and use it to browse error and warning messages.

     See Debugging components in the Application Server Toolkit.
  4. With WebSphere Process Server, run the dumpNameSpace command and pipe, redirect, or "more" the output so that it can be easily viewed.

     This command results in a display of all objects in WebSphere Process Server namespace, including the directory path and object name.
  5. If the object a client needs to access does not appear, use the administrative console to verify the following conditions.
     - The server hosting the target resource is started.

- The Web module or Enterprise JavaBean container hosting the target resource is running.
- The JNDI name of the target resource is properly specified.

If none of these steps solves the problem, see Troubleshooting and support for additional troubleshooting resources, including information about how to contact IBM Support.

- During the migration process, problems might occur while you are using the WBIPreUpgrade command or the WBIPostUpgrade command.
  - Problems can occur when you are using the WBIPreUpgrade command.
    - A ″Not found″ or ″No such file or directory″ message is returned.

      This problem can occur if you are trying to run the WBIPreUpgrade command from a directory other than the WebSphere Process Server version 6.1 *install_root*/bin directory. Verify that the WBIPreUpgrade script resides in the version 6.1 *install_root*/bin directory, and launch the file from that location.
    - The DB2® JDBC driver and DB2 JDBC driver (XA) cannot be found in the drop-down list of supported JDBC providers in the administrative console.

      The administrative console no longer displays deprecated JDBC provider names. The new JDBC provider names used in the administrative console are more descriptive and less confusing. The new providers will differ only by name from the deprecated ones.

      The deprecated names will continue to exist in the jdbc-resource-provider-templates.xml file for migration reasons (for existing JACL scripts for example); however, you are encouraged to use the new JDBC provider names in your JACL scripts.
    - You receive the following message:

      ```
      MIGR0108E: The specified WebSphere directory does not contain a WebSphere
      version that can be upgraded.
      ```

      This can occur if an incorrect directory was specified with the WBIPreUpgrade command.

    See the WBIPreUpgrade command.
  - Problems can occur when you are using the WBIPostUpgrade command.
    - A ″Not found″ or ″No such file or directory″ message is returned.

      This problem can occur if you are trying to run the WBIPostUpgrade command from a directory other than the WebSphere Process Server version 6.1 *install_root*\bin. Verify that the WBIPostUpgrade script resides in the version 6.1*install_root*\bin directory, and launch the file from that location.
    - When you migrate the federated nodes in a cell, you receive the following error messages:

      ```
      MIGR0304I: The previous WebSphere environment is being restored.
       com.ibm.websphere.management.exception.RepositoryException:
       com.ibm.websphere.management.exception.ConnectorException: ADMC0009E:
         The system failed to make the SOAP RPC call: invoke
      MIGR0286E: The migration failed to complete.
      ```

      A connection timeout occurs when the federated node tries to retrieve configuration updates from the deployment manager during the WBIPostUpgrade migration step for the federated node. Copying the entire configuration might take more than the connection timeout if the configuration that you are migrating to version 6.1 contains any of the following elements:

      - Many small applications

- A few large applications
- One very large application

If this occurs, modify the timeout value before running the WBIPostUpgrade command to migrate a federated node.

1. Go to the following location in the version 6.1 directory for the profile to which you are migrating your federated node:

   profile_root/properties

2. Open the soap.client.props file in that directory and find the value for the com.ibm.SOAP.requestTimeout property. This is the timeout value in seconds. The default value is 180 seconds.
3. Change the value of com.ibm.SOAP.requestTimeout to make it large enough to migrate your configuration. For example, the following entry would give you a timeout value of a half of an hour:

   ```
   com.ibm.SOAP.requestTimeout=1800
   ```

   Note: Select the smallest timeout value that will meet your needs. Be prepared to wait for at least three times the timeout that you select—once to download files to the backup directory, once to upload the migrated files to the deployment manager, and once to synchronize the deployment manager with the migrated node agent.
4. Go to the following location in the backup directory that was created by the WBIPreUpgrade command:

   migration_backup_directory/profiles/default/properties

5. Open the soap.client.props file in that directory and find the value for the com.ibm.SOAP.requestTimeout property.
6. Change the value of com.ibm.SOAP.requestTimeout to the same value that you used in the version 6.1 file.

- You receive the "Unable to copy document to temp file" error message. Here is an example:

  ```
  MIGR0304I: The previous WebSphere environment is being restored.
  com.ibm.websphere.management.exception.DocumentIOException: Unable to copy
  document to temp file:
    cells/sunblade1Network/applications/LARGEApp.ear/LARGEApp.ear
  ```

  Your file system might be full. If your file system is full, clear some space and rerun the WBIPostUpgrade command.

- You receive the following message:

  ```
  MIGR0108E: The specified WebSphere directory does not contain a WebSphere
  version that can be upgraded.
  ```

  The following possible reasons for this error exist:

  - An incorrect directory might have been specified when launching the WBIPreUpgrade command or the WBIPostUpgrade .
  - The WBIPreUpgrade command was not run.

- You receive the following error message:

  MIGR0253E: The backup directory *migration_backup_directory* does not exist.

  The following possible reasons for this error exist:

  - The WBIPreUpgrade command was not run before the WBIPostUpgrade command.
    1. Check to see if the backup directory specified in the error message exists.
    2. If not, run the WBIPreUpgrade command.

       See WBIPreUpgrade command.

3. Retry the WBIPostUpgrade command.

- An incorrect backup directory might be specified.

  For example, the directory might have been a subdirectory of the version 6.0.x tree that was deleted after the WBIPreUpgrade command was run and the older version of the product was uninstalled but before the WBIPostUpgrade command was run.

  1. Determine whether or not the full directory structure specified in the error message exists.
  2. If possible, rerun the WBIPreUpgrade command, specifying the correct full migration backup directory.
  3. If the backup directory does not exist and the older version it came from is gone, rebuild the older version from a backup repository or XML configuration file.
  4. Rerun the WBIPreUpgrade command.

- You decide that you need to run WBIPreUpgrade again after you have already run the WBIPostUpgrade command.

  During the course of a deployment manager or a managed node migration, WBIPostUpgrade might disable the old environment. If after running WBIPostUpgrade you want to run WBIPreUpgrade again against the old installation, you must run the migrationDisablementReversal.jacl script located in the old *install_root*/bin directory. After running this JACL script, your version 6.0.x environment will be in a valid state again, allowing you to run WBIPreUpgrade to produce valid results.

  For more information on scripting, see Getting started with scripting. Scripting, as described there, is available for WebSphere Process Server.

- A federated migration fails with message MIGR0405E.

  The migration that has taken place on your deployment manager as part of your federated migration has failed. For a more detailed reason for why this error has occurred, open the folder *your_node_name*_migration_temp located on your deployment manager node under the ...DeploymentManagerProfile/temp directory. For example:

  ```
  /websphere61/procserver/profiles/dm_profile/temp/nodeX
  _migration_temp
  ```

  The logs and everything else involved in the migration for this node on the deployment manager node are located in this folder. This folder will also be required for IBM support related to this scenario.

- WebSphere Process Server version 6.1 applications are lost during migration.

  If any of the version 6.1 applications fail to install during a federated migration, they will be lost during the synchronizing of the configurations. The reason that this happens is that one of the final steps of WBIPostUpgrade is to run a syncNode command. This has the result of downloading the configuration on the deployment manager node and overwriting the configuration on the federated node. If the applications fail to install, they will not be in the configuration located on the deployment manager node. To resolve this issue, manually install the applications after migration. If they are standard version 6.1 applications, they will be located in the *install_root*/installableApps directory.

  To manually install an application that was lost during migration, use the wsadmin command to run the install_*application_name*.jacl script that the migration tools created in the backup directory.

Linux In a Linux environment, for example, use the following parameters:

```
./wsadmin.sh -f migration_backup_directory/install_application_name.jacl
-conntype NONE
```

See the Wsadmin tool.

- WebSphere Process Server version 6.1 applications fail to install.

  Manually install the applications using the wsadmin command after WBIPostUpgrade has completed.

  To manually install an application that failed to install during migration, use the wsadmin command to run the install_*application_name*.jacl script that the migration tools created in the backup directory.

  Linux In a Linux environment, for example, use the following parameters:

  ```
  ./wsadmin.sh -f migration_backup_directory/install_application_name.jacl
  -conntype NONE
  ```

  See the Wsadmin tool, or see the WBIPostUpgrade command.

- Solaris When you use the migration wizard to migrate a profile from WebSphere Process Server version 6.0.x to version 6.1.x on a Solaris x64 processor-based system, the migration might fail during the WBIPostUpgrade step.

  You might see messages similar to the following in *profile_root*/logs/WASPostUpgrade.*time_stamp*.log:

  ```
  MIGR0327E: A failure occurred with stopNode.
  MIGR0272E: The migration function cannot complete the command.
  ```

  WebSphere Process Server version 6.0.x uses a Java virtual machine (JVM) in 32-bit mode. The migration wizard for WebSphere Process Server version 6.1.x calls the WBIPostUpgrade.sh script, which attempts to run the JVM for version 6.0.x in the 64-bit mode when the server stops the version 6.0.x node.

  Complete the following actions to remove the incomplete profile and enable WebSphere Process Server to correctly migrate the version 6.0.x profile:

  1. On a command line, change to the *install_root*/bin directory.

     For example, type the following command:

     ```
     cd /opt/IBM/WebSphere/Procserver/bin
     ```

  2. Locate the WBIPostUpgrade.sh script in the *install_root*/bin directory, and make a backup copy.
  3. Open the WBIPostUpgrade.sh or WBIPostUpgrade.bat file in an editor, and perform the following actions:
     a. Locate the following line of code:

        UNIX Linux

        ```
        "$binDir" /setupCmdLine.sh
        ```

        Windows

        ```
        call "%~dp0setupCmdLine.bat" %*
        ```

     b. Insert the following line of code after the code that was identified in the previous step:

        ```
        JVM_EXTRA_CMD_ARGS=""
        ```

     c. Save the changes.
  4. Repeat steps 2 through 4 with theWASPostUpgrade.sh or the WASPostUpgrade.bat file.

5. Use the following command to delete the incomplete version 6.1.x profile that was created during the migration process:

   ```
   install_root/bin/manageprofiles.sh -delete -profileName profile_name
   ```

6. Delete the *profile_root* directory of the version 6.1.x profile that was removed in the previous step.
7. Rerun the migration wizard.

- If you select the option for the migration process to install the enterprise applications that exist in the version 6.0.x configuration into the new version 6.1 configuration, you might encounter some error messages during the application-installation phase of migration.

  The applications that exist in the version 6.0.x configuration might have incorrect deployment information—usually, incorrect XML documents that were not validated sufficiently in previous WebSphere Process Server runtimes. The runtime now has an improved application-installation validation process and will fail to install these malformed EAR files. This results in a failure during the application-installation phase of WBIPostUpgrade and produces an ″E:″ error message. This is considered a ″fatal″ migration error.

  If migration fails in this way during application installation, you can do one of the following:

  – Fix the problems in the version 6.0.x applications, and then remigrate.
  – Proceed with the migration and ignore these errors.

    In this case, the migration process does not install the failing applications but does complete all of the other migration steps.

    Later, you can fix the problems in the applications and then manually install them in the new version 6.1 configuration using the administrative console or an install script.

- After migrating to a version 6.1 cell that contains or interoperates with version 6.0.x nodes that are not at WebSphere Process Server version 6.0.1.3 or later, the cluster function might fail.

  When starting these version 6.0.x servers, you might see the following problems:

  – You might see a first failure data capture (FFDC) log that shows a ClassNotFoundException error message. This exception is thrown from the RuleEtiquette.runRules method and looks something like the following example:

    ```
    Exception = java.lang.ClassNotFoundException
    Source = com.ibm.ws.cluster.selection.SelectionAdvisor.<init>
    probeid = 133
    Stack Dump = java.lang.ClassNotFoundException: rule.local.server
    at java.net.URLClassLoader.findClass(URLClassLoader.java(Compiled Code))
    at com.ibm.ws.bootstrap.ExtClassLoader.findClass(ExtClassLoader.java:106)
    at java.lang.ClassLoader.loadClass(ClassLoader.java(Compiled Code))
    at java.lang.ClassLoader.loadClass(ClassLoader.java(Compiled Code))
    at java.lang.Class.forName1(Native Method)
    at java.lang.Class.forName(Class.java(Compiled Code))
    at com.ibm.ws.cluster.selection.rule.RuleEtiquette.runRules(RuleEtiquette.java
    :154)at com.ibm.ws.cluster.selection.SelectionAdvisor.handleNotification
    (SelectionAdvisor.java:153)
    at com.ibm.websphere.cluster.topography.DescriptionFactory$Notifier.run
    (DescriptionFactory.java:257)
    at com.ibm.ws.util.ThreadPool$Worker.run(ThreadPool.java:1462)
    ```

  – You might see a java.io.IOException that looks something like the following example:

    ```
    Exception = java.io.IOException
    Source = com.ibm.ws.cluster.topography.DescriptionManagerA. update probeid
       = 362
    ```

```
Stack Dump = java.io.IOException
at com.ibm.ws.cluster.topography.ClusterDescriptionImpl.importFromStream
(ClusterDescriptionImpl.java:916)
at com.ibm.ws.cluster.topography.DescriptionManagerA.update
(DescriptionManagerA.java:360)
Caused by: java.io.EOFException
at java.io.DataInputStream.readFully(DataInputStream.java(Compiled Code))
at java.io.DataInputStream.readUTF(DataInputStream.java(Compiled Code))
at com.ibm.ws.cluster.topography.KeyRepositoryImpl.importFromStream
(KeyRepositoryImpl.java:193)
```

During migration, version 6.1 cluster information is distributed throughout the cell. WebSphere Process Server version version 6.0.x nodes that are not at version 6.0.1.3 or later fail to read this information.

To avoid this problem, upgrade all version 6.0.x nodes that will be contained in or interoperating with a version 6.1 cell to version 6.0.1.3 or later before migrating your deployment managers to version 6.1.

- After you migrate a managed node to version 6.1, the application server might not start.

  When you try to start the application server, you might see errors similar to those in the following example:

```
[5/11/06 15:41:23:190 CDT] 0000000a SystemErr  R
    com.ibm.ws.exception.RuntimeError:
com.ibm.ws.exception.RuntimeError:  org.omg.CORBA.INTERNAL:
   CREATE_LISTENER_FAILED_4
vmcid: 0x49421000  minor code: 56  completed: No
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr  R  at
com.ibm.ws.runtime.WsServerImpl.bootServerContainer(WsServerImpl.java:198)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr  R  at
com.ibm.ws.runtime.WsServerImpl.start(WsServerImpl.java:139)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr  R  at
com.ibm.ws.runtime.WsServerImpl.main(WsServerImpl.java:460)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr  R  at
com.ibm.ws.runtime.WsServer.main(WsServer.java:59)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr  R  at
sun.reflect.NativeMethodAccessorImpl.invoke0(Native Method)
[5/11/06 15:41:23:196 CDT] 0000000a SystemErr  R  at
sun.reflect.NativeMethodAccessorImpl.invoke(NativeMethodAccessorImpl.java:64)
[5/11/06 15:41:23:197 CDT] 0000000a SystemErr  R  at
sun.reflect.DelegatingMethodAccessorImpl.invoke
    (DelegatingMethodAccessorImpl.java:43)
```

  Change the port number at which the managed node's server is listening. If the deployment manager is listening at port 9101 for ORB_LISTENER_ADDRESS, for example, the server of the managed node should not be listening at port 9101 for its ORB_LISTENER_ADDRESS. To resolve the problem in this example, perform the following steps:

  1. On the administrative console, click **Application servers** → *server_name* → **Ports** → **ORB_LISTENER_ADDRESS**.
  2. Change the ORB_LISTENER_ADDRESS port number to one that is not used.

- In a network deployment environment, if the error `SRVE0026E: [Servlet Error]-[com/ibm/wbiservers/brules/BusinessRuleManager]: java.lang.NoClassDefFoundError` occurs when you access the Business Rules Manager after migrating, you must manually install the Business Rules Manager application on the deployment target before continuing with normal migration of that node. See Migrating Business Rules Manager in a network deployment environment for more information.
- If synchronization fails when you migrate a managed node to version 6.1, the server might not start.

You might receive messages similar to the following when you migrate a managed node to version 6.1:

```
ADMU0016I: Synchronizing configuration between node and cell.
ADMU0111E: Program exiting with error:
           com.ibm.websphere.management.exception.AdminException: ADMU0005E:
           Error synchronizing repositories
ADMU0211I: Error details may be seen in the file:
           /opt/WebSphere/61AppServer/profiles/AppSrv02/logs/syncNode.log
MIGR0350W: Synchronization with the deployment manager using the SOAP protocol
 failed.
MIGR0307I: The restoration of the previous WebSphere Application Server
 environment is complete.
MIGR0271W: Migration completed successfully, with one or more warnings.
```

These messages indicate the following:

- Your deployment manager is at a version 6.1 configuration level.
- The managed node that you are trying to migrate is at a version 6.1 configuration level on the deployment manager's repository (including applications).
- The managed node itself is not quite complete given that you did not complete the syncNode operation.

Perform the following actions to resolve this issue:

1. Rerun the syncNode command on the node to synchronize it with the deployment manager.

   See the syncNode command .
2. Run the GenPluginCfg command.

   See the GenPluginCfg command .

- In a network deployment environment, if the error `SRVE0026E: [Servlet Error]-[com/ibm/wbiservers/brules/BusinessRuleManager]: java.lang.NoClassDefFoundError` occurs when you access the Business Rules Manager after migrating, you must manually install the Business Rules Manager application on the deployment target before continuing with normal migration of that node. See Migrating Business Rules Manager in a network deployment environment for more information.

What to do next

If you did not find your problem listed, contact IBM support.

Troubleshooting migration from WebSphere InterChange Server

Find solutions to problems you encounter with migration as well as instructions for turning on logging and tracing.

Enabling logging and tracing for supported WebSphere InterChange Server APIs

Enable logging and tracing for the supported WebSphere InterChange Server APIs through the administrative console.

About this task

If your migrated application includes any supported WebSphere InterChange Server APIs, you can enable logging and tracing for them for troubleshooting purposes.

Procedure

1. Launch the administrative console.
2. From the left (navigation) panel, select **Troubleshooting > Logs and Trace**.
3. In the right panel, select the name of the server on which you want to enable logging and tracing.
4. In the right panel, under "General properties," select **Change Log Level Details**.
5. Select the Runtime tab. (Selecting the Runtime tab allows you to make this change in real time without requiring you to restart the server.)
6. Add the name of the package followed by `=all` to the list of logged packages in the box on the screen. Separate this new entry from any existing entries with a colon. For example, `CxCommon=all`. In this case, `CxCommon` is the name of the package for a set of supported WebSphere InterChange Server APIs. Specifying `all` enables all logging and tracing. See Supported WebSphere InterChange Server APIs for a list of the APIs, including their package names.
7. Select **Apply**.
8. To keep this configuration after the server is restarted, select the **Save runtime changes to configuration as well** check box.
9. Select **OK**.
10. When the next screen appears, select **Save** to save your changes.

Failure trying to serialize an object that is not serializable in a migrated BPEL file

If a serialization failure occurs in a BPEL file generated by the migration, you might be able to modify it to prevent the failure from occurring.

Problem: A serialization failure occurs in a custom snippet node of a business process execution language (BPEL) file generated by migration because an attempt is made to serialize an object that is not serializable.

Cause: In WebSphere InterChange Server, a Collaboration Template is compiled into a single Java class. In WebSphere Process Server, each node in a BPEL file might compile into a separate Java class. In WebSphere InterChange Server, a variable can be declared once and shared throughout the various steps of a Collaboration Template. To simulate that behavior in the migrated BPEL file, each variable used in a code snippet must be retrieved at the start of the snippet and saved at the end of the snippet. Variables defined in WebSphere InterChange Server Port definitions become BPEL variables. These are retrieved into BusObj variables at the beginning of each snippet (if referenced in the snippet) and saved back to the BPEL variables at the end of each snippet. For example, a retrieval at the beginning of snippets looks like this:

```
BusObj tempBusObj = null;if (tempBusObj_var != null) { tempBusObj =
      new BusObj(tempBusObj_var); };
```

and a save at the end of snippets looks like this:

```
if (tempBusObj == null) { tempBusObj_var = null; } else { tempBusObj_var =
      tempBusObj.getBusinessGraph(); }
```

Other variables used in the WebSphere InterChange Server snippet code are serialized and stored as a String in a BPEL variable named *CollabTemplateName_var*. These variables are deserialized at the beginning of each BPEL snippet, and then serialized and saved at the end of each BPEL Snippet that they are referenced in. For example, objects are retrieved like this:

```
BusObj tempBusObj = (BusObj)BaseCollaboration.deserialize
  (FrontEndCollab_var.getString("tempBusObj"));
```

and objects are saved like this:

```
FrontEndCollab_var.setString("tempBusObj", BaseCollaboration.serialize(tempBusObj));
```

If the type of the object being serialized is not serializable, then using serialize and deserialize will fail when the BPEL is run.

Solution: After migration, modify the BPEL file as follows:

- For any variable that is not Java-serializable, update the BPEL snippets to remove the serialization and deserialization statements. If the variable needs to be shared across snippets (instead of being re-created in each snippet) another method must be used to preserve the value of the variable across snippets.
- Manually define BPEL variables for variables of type BusObj that are not declared in the WebSphere InterChange Server Port definitions but are used on Partner Invokes. This is a manual step because variables used during invokes in WebSphere Process Server must be strongly typed, and the migration tools cannot accurately determine that type from the WebSphere InterChange Server snippets.

  Note: The naming convention used by the migration tools is to add `_var` to the name of the variable in the snippet code when naming the BPEL variables. For example, for a variable called tempBusObj in the snippet code, the migration tools will create a BPEL variable named `tempBusObj_var`.
- For variables that must be declared manually as BPEL variables, change the BPEL snippet code so that it uses the ″deserialize/serialize″ method of preserving these variables rather than the ″retrieve from/store into BPEL variable″ method of preserving these variables.

Chapter 7. 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 J2EE 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 J2C activation specifications created prior to 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 J2C activation specifications (see Deleting SIBus destinations and Deleting J2C 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 39.

   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 J2C activation specifications

The system builds J2C 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 J2C 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 J2C 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

SIBus destinations are the connections that make services available to applications. There will be times that you will have to remove destinations.

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 J2C 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.
   Navigate to the panel by clicking **Service integration > Buses**
3. Select the SCA system bus destinations.
   In the display, click on **SCA.SYSTEM.*cellname*.Bus**, where *cellname* is the name of the cell that contains the module with the destinations you are deleting.
4. Delete the destinations that contain a module name that matches the module that you are removing.
   a. Click on the check box next to the pertinent destinations.
   b. Click **Delete.**

Results

The panel displays only the remaining destinations.

What to do next

Delete the J2C activation specifications related to the module that created these destinations.

Chapter 8. Troubleshooting WebSphere Process Server administration

Troubleshooting is the process of finding and eliminating the cause of a problem. This group of topics helps you identify and resolve problems that can occur during typical administration tasks.

For information on troubleshooting Business Process Choreographer or Common Event Infrastructure components, see one of the following locations:

- The WebSphere Process Server for Multiplatforms, version 6.1, information center
- The *Business Process Choreographer* PDF
- The *Common Event Infrastructure* PDF

Using cross-component tracing for applications

Cross-component tracing allows you to identify trace.log data that is associated with WebSphere Process Server and WebSphere Enterprise Service Bus modules and components. The trace.log data can include error and event information, such as corrupted data or runtime exceptions, captured during Service Component Architecture (SCA) processing. The input and output data passing between WebSphere Process Server and WebSphere Enterprise Service Bus components can also be captured and used for problem determination using WebSphere Integration Developer.

Enabling Cross-Component Trace

Enable cross-component tracing to collect error and event information associated with WebSphere Process Server and WebSphere Enterprise Service Bus modules and components captured during processing.

Before you begin

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

About this task

To enable cross-component tracing, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Troubleshooting** → **Cross-Component Trace** to display the Cross-Component Trace page.
2. Select a server for which you want to enable tracing. Note the status of the selected servers in the **Status** column: **running** or **not running**. If the status is **not running**, runtime tracing for that server is disabled and you can specify only the configuration value. In this case, the trace level takes effect only when the server starts or restarts.
3. In the **Configuration** or **Runtime** columns, select **enable** from the dropdown box for each server for which you want to enable tracing. Enable tracing in **Configuration** to collect data when the server starts or restarts. Enable tracing in **Runtime** to collect data when the server is currently running.

4. After you have specified the settings, click **OK** to save the settings.

Results

Collected data is added to the trace.log file and is purged as those files are purged. See "Troubleshooting Service Component Architecture (SCA) processing and call chains" on page 46 for more information.

Enabling Cross-Component Tracing with data snapshot

Enable cross-component tracing with data snapshot to collect data associated with Service Component Architecture (SCA) processing and call chain data associated with WebSphere Process Server and WebSphere Enterprise Service Bus modules and components.

Before you begin

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

About this task

To enable cross-component tracing, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Troubleshooting** → **Cross-Component Trace** to display the Cross-Component Trace page.
2. Select the server for which you want to enable tracing. Note the status of the selected servers in the **Status** column: **running** or **not running**. If the status is **not running**, runtime tracing for that server is disabled and you can specify only the configuration value. In this case, the trace level takes effect only when the server starts or restarts.
3. In the **Configuration** or **Runtime** columns, select **enable with data snapshot** from the dropdown box for each server for which you want to enable tracing. Enable tracing in **Configuration** to collect data when the server starts or restarts. Enable tracing in **Runtime** to collect data when the server is currently running.
4. After you have specified the settings, click **OK** to save the settings.

Results

Collected SCA data is added to the systemout.log and trace.log files and is purged as those files are purged. Input and output data passing between WebSphere Process Server and WebSphere Enterprise Service Bus components is captured and additional files are created in the logs\XCT directory. This data can be used for problem determination by WebSphere Integration Developer for problem determination. Deleting these files when they are no longer needed is a task for the administrator. See "Troubleshooting Service Component Architecture (SCA) processing and call chains" on page 46 for more information.

Disabling Cross-Component Trace

Disable cross-component tracing to stop the collection of error and event information associated with WebSphere Process Server and WebSphere Enterprise Service Bus modules and components captured during Service Component Architecture (SCA) processing.

Before you begin

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

About this task

To disable cross-component tracing, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Troubleshooting** → **Cross-Component Trace** to display the Cross-Component Trace page.
2. Select a server for which you want to disable tracing. Note the status of the selected servers in the **Status** column: **running** or **not running**. If the status is **not running**, runtime tracing for that server is disabled and you can specify only the configuration value. In this case, the disabling tracing takes effect only when the server starts or restarts.
3. In the **Configuration** or **Runtime** columns, select **disable** from the dropdown box for each server for which you want to disable tracing.
4. After you have specified the settings, click **OK** to save the settings.

Results

Cross-component tracing is disabled for the selected servers. No data is collected.

Deleting data collected by Cross-Component Trace

When Cross-Component trace adds data to log files, and if enabled with data capture, additional files are created in the logs\XCT directory. The data added to the systemout.log and trace.log files does not need to be deleted as these files are automatically deleted by WebSphere Application Server. The data capture files need to be deleted from the logs\XCT directory manually when they are no longer needed. Delete these files after using WebSphere Integration Developer for problem determination or after the log files that refer to the files have been deleted by WebSphere Application Server.

Before you begin

You must have read and write access to the logs directories of each server.

About this task

To delete data collected by cross-component tracing, use the following procedure.

Procedure

1. Go to the logs\XCT directory in which the data was captured and move the contents of the directory to a location where it can be viewed by WebSphere Integration Developer for problem determination.
2. If you determine that the captured data is not needed for problem determination, then manually delete the contents of the logs\XCT directory.

Results

The captured data is deleted.

Troubleshooting Service Component Architecture (SCA) processing and call chains

Cross-component tracing 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 trace.log data can include error and event information, such as corrupted data or runtime exceptions, captured during SCA processing. The input and output data passing between WebSphere Process Server and WebSphere Enterprise Service Bus components can also be captured and 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**. On this page, you can select the servers from which you want to collect trace data. Use the **Configuration** column to specify trace settings for servers that will be used when the server starts or restarts. Use the **Runtime** column to specify trace settings for servers that are running.

For each server, you can use any of the following settings:

enable
: This setting enables tracing for SCA processing. Data collected from this setting is added to the systemout.log and trace.log files and is purged as those files are purged.

enable with data snapshot
: This setting enables tracing for SCA processing and input and output data that passes between WebSphere Process Server and WebSphere Enterprise Service Bus components. Data from SCA processing is added to the systemout.log and trace.log files. Input and output data from WebSphere Process Server and WebSphere Enterprise Service Bus components is placed in files that are created in the logs\XCT directory.

disable
: This setting disables tracing on the selected server.

Handling and deleting collected data

- Data collected from SCA processing is added to added to the systemout.log and trace.log files and is purged as those files are purged.
- Input and output data added as files in the logs\XCT directory can be moved to a location where it can be viewed by WebSphere Integration Developer for problem determination. The files can then be manually deleted. These input and output files in logs\XCT are related to the systemout.log and trace.log files that were created at the same time. When WebSphere Application Server deletes old systemout.log and trace.log files, the associated input and output files in logs\XCT can also be deleted. Generally, there will be many input and output files for a given systemout.log file. The timestamps in the systemout.log and trace.log files can be used to identify what input and output files to delete. It is safe to delete all input and output files older than the oldest date in the systemout.log and trace.log files.

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 with reduced performance during an advanced search	"Advanced search feature is not optimized"
I am having trouble entering values in the Search page's **By Date** tab	"Values in the By Date tab 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 48
I am having trouble with failed events not being created	"Failed events are not being created" on page 48

Advanced search feature is not optimized

The failed event manager's advanced search feature is not optimized. Therefore, you may experience reduced performance when using the Advanced search tab with a large set of failed events.

Values in the By Date tab automatically change to default if entered incorrectly

The Search page's **By Date** tab contains two fields: **From Date** and **To Date**. Both fields are required. The values are locale-dependendent, and they must be formatted exactly as shown in the example above the field. 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 machine 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 actually 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 Service Component Architecture and WebSphere MQ communications

Communication between Service Component Architecture (SCA) modules and WebSphere MQ queue managers depends on the binding between the imports and exports within the SCA module and the queues in WebSphere MQ servers. Use this information to determine the servers that are not processing WebSphere MQ messages.

Before you begin

This task assumes that you have noticed requests dependant on WebSphere MQ are not being processed and that you have access to the administrative console. You should also either have the ability to make changes to the WebSphere MQ queue manager or be in contact with the WebSphere MQ administrator.

About this task

Service Component Architecture (SCA) modules depend on the bindings between the server and the WebSphere MQ queue manager. Communications between the two entities could keep messages from processing completely. The following steps should help you discover the cause of the disruption and what to do to get the messages processed again.

Procedure

1. Display the SCA module communicating with WebSphere MQ to make sure it is still processing. Navigate to this page using **Applications > SCA Modules**.

2. Display the queue manager to make sure it is still operational. Perform this task at the WebSphere MQ administration console.
3. Display the bindings between the SCA module and the queue manager to make sure the binding is correct. If the binding is incorrect, change the binding. Navigate to this page using **Applications** → **SCA modules** >*moduleName* > **Imports | Exports** > *importName | exportName* > **Bindings** > *bindingName* [type].
4. Locate any messages that may indicate failed transactions. You will have to investigate system, SCA-specific message areas, WebSphere MQ-specific message areas, the failed event queue and other locations to determine what has failed.
   a. Examine SystemOut.log for any messages that would indicate processing failures.

      If there is an WebSphere MQ error, there will be an MQException linked somewhere in the stack trace with a WebSphere MQ reason code (for example, 2059 is "queue manager unavailable").
   b. Check AMQERR*xx*.LOG and the WebSphere MQ FFDC files to determine the cause of a WebSphere MQ error.
   c. Examine the application queues to determine if there are any unprocessed messages. Make sure you examine both WebSphere MQ and Service Integration Bus (SIB) queues.
   d. Examine the WebSphere MQ dead letter queue and the SIB exception destination.
   e. Examine the failed event queue to determine if there are any messages related to the applications of interest. See Finding failed events for information about locating the failed events. See "Managing WebSphere Process Server failed events" for information about locating the failed events.

Troubleshooting event sequencing

Refer to the information in this topic if you are experiencing difficulty with event sequencing.

Problems with the event sequencing qualifier

Ensure that your component definition is correct:

- Is the event sequencing qualifier set on the method? Event sequencing validation fails if the qualifier is erroneously set on the interface.
- Is the parameter name valid?
- Is the xpath element valid, and does it correctly resolve to a primitive?
- Is there a single eventSequencing element for the method? Each method supports only one eventSequencing element.
- Is there a single keySpecification element for the method? Each method supports only one keySpecification element.

Deadlocks

Deadlocks occur when an invoked operation with a lock invokes another operation on the same component using the same event sequencing key and group. You can resolve a deadlock by using the esAdmin command to list and release the current lock.

To avoid deadlocks, carefully consider dependencies when implementing event sequencing. Ensure that operations with circular dependencies are in different event sequencing groups.

Deadlocks with a BPEL process

Deadlocks can occur when event sequencing is used with Business Process Execution Language (BPEL) processes. Deadlocks are caused by setting event sequencing qualifiers on operations that correspond to both of the following activities:

- Multiple instantiating receive or pick activities, where the `createInstance` attribute is set to `yes`
- Correlation set specifications with an `initiation` attribute set to `join`

Resolve this type of deadlock by using the esAdmin command to list and release the current lock. To prevent further deadlocks, ensure that these types of dependent operations are put into different event sequencing groups.

Performance issues

If you are experiencing memory problems on the messaging engine server used for event sequencing components, try modifying the runtime event sequencing property maxActiveMessages in the *install_root*/properties/eventsequencing.properties file.

The maxActiveMessages property defines the number of messages currently locked on a component destination; too many large messages can negatively affect performance and cause memory problems. Note that a value of `0` (zero) means that an unlimited number of messages are allowed. By default, the maxActiveMessages property is set to 100.

To modify the maxActiveMessages property, perform the following steps.

1. Open the eventsequencing.properties file in a text editor.
2. Make the appropriate modifications for your environment.
3. Save and close the file.
4. Stop and restart any applications that are part of the event sequencing component in order for the changes to take effect.

Troubleshooting your deployment environment

When processing appears sluggish or requests fail, use a focused approach to determine the source of the problem in the environment. The approach described is for non-standalone server environments.

Before you begin

You must be logged into the administrative console of the deployment manager to perform this task.

Required security role for this task: When security and role-based authorization are enabled, you must log in to the administrative console as an administrator or operator to perform this task.

About this task

Investigate the state of your deployment environment if you notice any of the following symptoms:
- Unavailable applications
- Sluggish applications
- Stopped applications
- Decreased throughput
- Sluggish performance

Procedure

1. Display the topology layout that describes this deployment environment to determine the status of the topology.
2. Display the topology to determine the state of the various roles in the topology. Note the roles with unexpected states or warning for further investigation.
3. Locate the nodes that are causing the error state for each role.
4. Make sure all nodes are synchronized.

   On the Nodes page in the administrative console select any unsynchronized nodes and click **Synchronize**.
5. Make sure that the messaging engines associated with all the buses are running.

   If they are not running, stop and start the messaging engines.
6. Locate the logs associated with the nodes in error and view the logs for error messages.
7. Take any actions prescribed by the error messages to affect the correction.
8. Correct any errors and restart the affected nodes.

Results

The nodes previously in error start and the status of the topology becomes "running."

What to do next

Restart any affected applications

Troubleshooting the business rules manager

Some of the problems you may 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.

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 global 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.

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 global 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 messaging bindings

Various error conditions can occur with bindings that are specific to the type of binding.

About this task

The manner in which error conditions are handled depends upon the type of binding concerned.

Troubleshooting JMS and WebSphere MQ JMS bindings

You can diagnose and fix problems with JMS and WebSphere MQ JMS bindings.

Implementation exceptions

About this task

In response to various error conditions, the JMS and MQ JMS import and export implementation can return one of two types of exceptions:

- ServiceBusinessException: this exception is returned if exception specified on the service business interface (WSDL port type or Java interface) occurred.
- ServiceRuntimeException: raised in all other cases. In most cases, the cause exception will contain the original exception. In the J2C case that would be ResourceException and in the JMS case, it would be JMSException.

When an XML schema has a type defined without a global element, the JMS bindings (JMSDataBindingImplXML and JMSDataBindingImplJava) cannot resolve the type to an element.

Schema:

```
<?xml version="1.0" encoding="UTF-8"?>
<schema xmlns="http://www.w3.org/2001/XMLSchema"
targetNamespace="http://www.ibm.com" xmlns:tns="http://www.ibm.com">
 <!-- global element required but missing -->

<complexType name="Quote">
  <sequence>
   <element name="symbol" type="string"></element>
```

```
  <element name="price" type="float"></element>
  </sequence>
 </complexType>
</schema>
```

If you receive one of the following exceptions:

```
com.ibm.websphere.sca.ServiceRuntimeException:
caused by: java.lang.IllegalArgumentException:
{Quote}Quote is not corresponding to a global element.
```

Or

```
[8/25/06 10:20:40:938 PDT] 00000054 FFDC          Z
com.ibm.ws.sca.databinding.impl.DataBindingImplXML
com.ibm.ws.sca.databinding.impl.DataBindingImplXML#002 Exception:

org.eclipse.emf.ecore.xmi.FeatureNotFoundException:
Feature 'Quote' not found. (sca:/dataObject.xml, 2, 126)
```

This might mean you need to define a global element:

```
<?xml version="1.0" encoding="UTF-8"?>
<schema xmlns="http://www.w3.org/2001/XMLSchema"
targetNamespace="http://www.ibm.com" xmlns:tns="http://www.ibm.com">

<element name="Quote" type="tns:Quote"></element> <!-- global element required -->

<complexType name="Quote">
  <sequence>
   <element name="symbol" type="string"></element>
   <element name="price" type="float"></element>
  </sequence>
 </complexType>

</schema>
```

JMS-based SCA messages not appearing in the failed event manager

About this task

If SCA messages originated through a JMS interaction fail, you would expect to find these messages in the failed event manager. If such messages are not appearing in the failed event manager, ensure that the underlying SIB destination of the JMS Destination has a maximum failed deliveries value greater than 1. Setting this value to 2 or more enables interaction with the failed event manager during SCA invocations for the JMS bindings.

WebSphere MQ JMS-based SCA messages not appearing in the failed event manager

About this task

If SCA messages originated through a WebSphere MQ JMS interaction fail, you would expect to find these messages in the failed event manager. If such messages are not appearing in the failed event manager, ensure that the value of the maximum retries property on the underlying listener port is equal to or greater than 1. Setting this value to 1 or more enables interaction with the failed event manager during SCA invocations for the MQ JMS bindings.

Faults

About this task

The argument that is expected to be passed to the JMSDataBinding and JMSObjectBinding depends on the interface operation and the input, output and fault types.

For faults, the outDataBindingType specified on the method binding is used. If none is specified, the binding level dataBindingType is used for all serialization and deserialization.

If the fault type is simple, a string is set on the JMS DataBinding representing the fault message. In addition IsBusinessException is set to true.

If the fault type is a data object, then a data object is set on the JMS DataBinding which represents the fault message. This scenario requires the use of JMSDataBinding.

Messages containing faults are handled by the JMS data bindings. A Boolean header property IsBusinessException is intercepted by the data binding. If the value is true, the data binding informs the runtime that the payload contains fault data.

If you are working with a custom data binding, you need to take the following steps to handle faults correctly. Default implementations handle faults without user intervention.

Procedure

1. For JMS exports, use the setBusinessException (boolean isBusinessException) method on the JMSDataBinding interface to indicate that data object or object specified on the data binding is a fault object and the message created by the binding should be constructed accordingly. The Data Binding is then responsible to specify the isBusinessException appropriately.
2. For JMS imports use the isBusinessException() method on the JMSDataBinding interface to indicate whether the message contains a fault.

   The data binding gets the value of the header property indicating a fault defined in the payload. After the runtime passes the JMS message to the data binding, it invokes isBusinessException() on the data binding. If returned value is false, the message is processed normally, otherwise, the ServiceBusinessException is returned to the caller. The data object or object produced by the binding is set on a ServiceBusinessException and it is returned to the caller.

Misusage scenarios: comparison with WebSphere MQ bindings

About this task

The WebSphere MQ JMS binding is designed to interoperate with JMS applications deployed against WebSphere MQ, which exposes messages according to the JMS message model. The WebSphere MQ import and export, however, are principally designed to interoperate with native WebSphere MQ applications, and expose the full content of the WebSphere MQ message body to mediations.

The following scenarios should be built using the WebSphere MQ JMS binding, not the WebSphere MQ binding:

- Invoking a JMS message-driven bean (MDB) from an SCA module, where the MDB is deployed against the WebSphere MQ JMS provider. This would be implemented using a WebSphere MQ JMS import.

- Allowing the SCA module to be called from a J2EE component servlet or EJB via JMS. This would be implemented using a WebSphere MQ JMS export.
- Mediating the contents of a JMS MapMessage, in transit across WebSphere MQ. This would use WebSphere MQ JMS export and import in conjunction with the appropriate data binding.

There are situations in which the WebSphere MQ binding and WebSphere MQ JMS binding might be expected to interoperate. In particular, when you are bridging between J2EE and non-J2EE WebSphere MQ applications, use a WebSphere MQ export and WebSphere MQ JMS import (or vice versa) in conjunction with appropriate data bindings, mediation module, or both.

Troubleshooting generic JMS bindings

Certain failure conditions can occur with a generic JMS binding.

About this task

Various error conditions can occur with generic JMS bindings.

Troubleshooting generic JMS exceptions

In response to various error conditions, the generic JMS import and export implementation can return an exception.

About this task

In response to various error conditions, the generic JMS import and export implementation can return one of two types of exceptions:

- ServiceBusinessException – this exception is returned if the exception specified on the service business interface (WSDL port type or Java interface) occurred.
- ServiceRuntimeException – raised in all other cases. In most cases, the *cause* exception will contain the original exception. In the case of JMS, it would be JMSException.

Troubleshooting generic JMS message expiry

A request message by the JMS provider is subject to expiration.

About this task

Request expiry refers to the expiration of a request message by the JMS provider when the JMSExpiration time on the request message is reached. As with other JMS bindings, the generic JMS binding handles the request expiry by setting expiration on the callback message placed by the import to be the same as for the outgoing request. Notification of the expiration of the callback message will indicate that the request message has expired and the client should be notified by means of a business exception.

If the callback destination is moved to the third-party provider, however, this type of request expiry is not supported.

Response expiry refers to the expiration of a response message by the JMS provider when the JMSExpiration time on the response message is reached.

Response expiry for the generic JMS binding is not supported, because the exact expiry behavior of a third-party JMS provider is not defined. You can, however, check that the response is not expired if and when it is received.

For outbound request messages, the JMSExpiration value will be calculated from the time waited and from the requestExpiration values carried in the asyncHeader, if set.

Troubleshooting generic JMS connection factory errors

When you define certain types of Connection Factory in your generic JMS provider, you may receive an error message when you try to start an application. You can modify the external Connection Factory to avoid this problem.

About this task

When launching an application you may receive the following error message: "MDB Listener Port JMSConnectionFactory type does not match JMSDestination type"

This problem can arise when you are defining external connection factories. Specifically the exception can be thrown when you create a JMS 1.0.2 Topic Connection Factory, instead of a JMS 1.1 (unified) Connection Factory (i.e., one that is able to support both point-to-point and publish/subscribe communication).

To resolve this issue take the following steps:

Procedure

1. Access the generic JMS provider that you are using.
2. Replace the JMS 1.0.2 Topic Connection Factory that you defined with a JMS 1.1 (unified) Connection Factory.

Results

When you launch the application with the newly defined JMS 1.1 Connection Factory you should no longer receive an error message.

Troubleshooting WebSphere MQ bindings

You can diagnose and fix faults and failure conditions that occur with WebSphere MQ bindings.

About this task

The primary failure conditions of WebSphere MQ bindings are determined by transactional semantics, by WebSphere MQ configuration, or by reference to existing behavior in other components. The primary failure conditions include:

- Failure to connect to WebSphere MQ queue manager or queue: a failure to connect to WebSphere MQ, to receive messages, will result in the MDB ListenerPort failing to start. This condition will be logged in the WebSphere Application Server log. Persistent messages will remain on the WebSphere MQ queue until they are successfully retrieved (or expired by WebSphere MQ). A failure to connect to WebSphere MQ, to send outbound messages, will cause rollback of the transaction controlling the send.
- Failure to parse inbound message, or to construct outbound message: a failure in the data binding will cause rollback of the transaction controlling the work.
- Failure to send outbound message: a failure to send a message will cause rollback of the relevant transaction

- Multiple or unexpected response messages: the import expects only one response message for each request message; when a response has arrived, the record is deleted. If response messages arrive unexpectedly, they will be discarded as with the JMS import.

Misusage scenarios: comparison with WebSphere MQ JMS bindings

About this task

The WebSphere MQ import and export are principally designed to interoperate with native WebSphere MQ applications, and expose the full content of the WebSphere MQ message body to mediations. The WebSphere MQ JMS binding, however, is designed to interoperate with JMS applications deployed against WebSphere MQ, which exposes messages according to the JMS message model.

The following scenarios should be built using the WebSphere MQ JMS binding, not the WebSphere MQ binding:

- Invoking a JMS message-driven bean (MDB) from an SCA module, where the MDB is deployed against the WebSphere MQ JMS provider. This would be implemented using a WebSphere MQ JMS import.
- Allowing the SCA module to be called from a J2EE component servlet or EJB via JMS. This would be implemented using a WebSphere MQ JMS export.
- Mediating the contents of a JMS MapMessage, in transit across WebSphere MQ. This would use WebSphere MQ JMS export and import in conjunction with the appropriate data binding.

There are situations in which the WebSphere MQ binding and WebSphere MQ JMS binding might be expected to interoperate. In particular, when you are bridging between J2EE and non-J2EE WebSphere MQ applications, use a WebSphere MQ export and WebSphere MQ JMS import (or vice versa) in conjunction with appropriate data bindings and/or mediation module.

If WebSphere MQ cannot deliver a message to its intended destination, usually due to configuration errors, it will send messages instead to a nominated dead-letter queue. In doing so it adds a dead-letter header to the start of the message body; this header contains failure reasons, the original destination, and other information.

Profile-specific log files

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

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.

i5/OS **On i5/OS platforms:** The variable *user_data_root* represents the default user data directory.

For more information see Default installation directories for the product, profiles, and tools.

Table 2. 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 these directories: • Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs/ffdc • Windows **On Windows platforms:** *profile_root*\logs\ffdc • i5/OS **On i5/OS platforms:** *profile_root*/logs/ffdc	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: • **Deployment manager profile:** – Log file — *deployment_manager_name*_exception.log. – Text files — *deployment_manager_name_hex_id_date_time*.txt. • **Custom profile:** • – Log file(s) — *node_agent_name*_exception.log and *server_name*_exception.log . – Text files — *node_agent_name*(or)*server_name_hex_id_date_time*.txt. • **Stand-alone profile:** – Log file — *server_name*_exception.log. – Text files — *server_name_hex_id_date_time*.txt.
Deployment manager logs (deployment manager profiles only) are found in these directories: • Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs/*deployment_manager_name* • Windows **On Windows platforms:** *profile_root*\logs*deployment_manager_name* • i5/OS **On i5/OS platforms:** *profile_root*/logs/*deployment_manager_name*	You will work primarily with four log files in this directory: • 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. • SystemErr.log — Contains error and exception messages generated by the deployment manager during runtime. Continually updated while server is running. • SystemOut.log — Contains all messages, including error, warning, and information messages generated by the deployment manager during runtime. Continually updated while server is running.

Table 2. Profile-specific log files updated during runtime (continued)

Log	Contents
Node agent logs (custom profiles only) are found in these directories: • Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs/*node_agent_name* • Windows **On Windows platforms:** *profile_root*\logs*node_agent_name* • i5/OS **On i5/OS platforms:** *profile_root*/logs/*node_agent_name*	You will work primarily with four log files in this directory: • 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. • SystemErr.log — Contains error and exception messages generated by the node agent during runtime. Continually updated while node agent is running. • SystemOut.log — Contains all messages, including error, warning, and information messages generated by the node agent during runtime. Continually updated while the node agent is running.
Server logs (custom and stand-alone profiles only) are found in these directories: • Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs/*server_name* • Windows **On Windows platforms:** *profile_root*\logs*server_name* • i5/OS **On i5/OS platforms:** *profile_root*/logs/*server_name*	You will work primarily with four log files in this directory: • 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. • SystemErr.log — Contains error and exception messages generated by the server during runtime. Continually updated while server is running. • SystemOut.log — Contains all messages, including error, warning, and information messages generated by the server during runtime. Also contains any events being monitoring that are emitted from the Common Event Infrastructure (CEI), in Common Base Event format. These events may also include the level of business object data (FINE, FINER, or FINEST) that is specified for the monitor. Continually updated while the server is running.

Table 2. Profile-specific log files updated during runtime (continued)

Log	Contents
Node federation log files are found in these directories (only applies to non-deployment manager profiles): • Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs • Windows **On Windows platforms:** *profile_root*\logs • i5/OS **On i5/OS platforms:** *profile_root*/logs	Two log files are generated when you attempt to federate a custom, augmented, or stand-alone profile to a deployment manager: • 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.
The location of the Integrated Solutions Console application deployment log file is listed here (only for deployment manager and stand-alone profiles): • Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs/iscinstall.log • Windows **On Windows platforms:** *profile_root*\logs\iscinstall.log • i5/OS **On i5/OS platforms:** *profile_root*/logs/iscinstall.log	The iscinstall.log file contains information regarding the deployment of the administrative console application in a deployment manager or stand-alone profile.
The location of the Installation Verification Tool log file is listed here (only for deployment manager and stand-alone profiles): • Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs/ivtClient.log • Windows **On Windows platforms:** *profile_root*\logs\ivtClient.log • i5/OS **On i5/OS platforms:** *profile_root*/logs/ivtClient.log	This log file contains the output generated by the Installation Verification Tool. You can start this program from the First Steps console after you create a deployment manager or stand-alone profile. The log contains basic configuration information and the messages that are displayed when you run the tool.
The location of the log file detailing the commands generated for a profile creation is listed here: • Linux UNIX **On Linux and UNIX platforms:** *profile_root*/logs/updateserverpolicy.log • Windows **On Windows platforms:** *profile_root*\logs\updateserverpolicy.log • i5/OS **On i5/OS platforms:** *profile_root*/logs/updateserverpolicy.log	This file contains the sequence of commands used by the product to set server environment variables and create a profile. All profile types will contain this file.

Related tasks

Chapter 4, “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 comprise the second group.

Installation and profile creation log files

Various log files are created during installation and uninstallation of WebSphere Process Server and during profile creation, augmentation, and deletion. Consult the applicable logs if problems occur during these procedures.

Chapter 9. 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 Network Deployment, version 6.1.

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

Chapter 10. 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 Version 6.1 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, version 6.1. 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.

- 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

http://publib.boulder.ibm.com/infocenter/dmndhelp/v6r1mx/topic/com.ibm.websphere.wps.612.doc/doc/cmon_businessevents.html

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 (SCA) processing and call chains

Cross-component tracing 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 trace.log data can include error and event information, such as corrupted data or runtime exceptions, captured during SCA processing. The input and output data passing between WebSphere Process Server and WebSphere Enterprise Service Bus components can also be captured and 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**. On this page, you can select the servers from which you want to collect trace data. Use the **Configuration** column to specify trace settings for servers that will be used when the server starts or restarts. Use the **Runtime** column to specify trace settings for servers that are running.

For each server, you can use any of the following settings:

enable
: This setting enables tracing for SCA processing. Data collected from this setting is added to the systemout.log and trace.log files and is purged as those files are purged.

enable with data snapshot
: This setting enables tracing for SCA processing and input and output data that passes between WebSphere Process Server and WebSphere Enterprise Service Bus components. Data from SCA processing is added to the systemout.log and trace.log files. Input and output data from WebSphere Process Server and WebSphere Enterprise Service Bus components is placed in files that are created in the logs\XCT directory.

disable
: This setting disables tracing on the selected server.

Handling and deleting collected data

- Data collected from SCA processing is added to added to the systemout.log and trace.log files and is purged as those files are purged.
- Input and output data added as files in the logs\XCT directory can be moved to a location where it can be viewed by WebSphere Integration Developer for problem determination. The files can then be manually deleted. These input and

output files in logs\XCT are related to the systemout.log and trace.log files that were created at the same time. When WebSphere Application Server deletes old systemout.log and trace.log files, the associated input and output files in logs\XCT can also be deleted. Generally, there will be many input and output files for a given systemout.log file. The timestamps in the systemout.log and trace.log files can be used to identify what input and output files to delete. It is safe to delete all input and output files older than the oldest date in the systemout.log and trace.log files.

Managing failed events

The WebSphere Process Server Recovery service monitors for failed operations between Service Component Architecture (SCA) components. If an operation fails, the Recovery service captures data about the event and the failure. You can then use the failed event manager to view, modify, resubmit, or delete the failed event.

What is a failed event?

In the context of WebSphere Process Server, an event is a request that is received by a WebSphere Process Server application. It can come from an external source (such as an inbound application adapter) or an external invocation to a Web service. The event is comprised 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 WebSphere Process Server 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 WebSphere Process Server Recovery service handles the following types of failed events:

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

The Recovery service does not handle failures from synchronous invocations or from Business Process Execution Language (BPEL) asynchronous request/reply invocations.

Failed 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 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.

How are failed events managed?

The Recovery service sends failed 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.

An administrator uses the failed event manager to browse and manage all WebSphere Process Server 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**.

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

Before you begin

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 on the server, with references to topics for conducting other searches based on session ID, source, destination, date, business object type, exception text, or a combination of those criteria.

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. Click **Failed events on this server → 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 modify data in a failed event, resubmit it, or delete it.

Searching for failed events by session ID

Use the Search page's **By Session** tab to find only those events that failed within a specific session.

Before you begin

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

About this task

Every event executes within a session; if that event fails, the failed event manager encapsulates specific session information for the failed execution branch in the Session ID parameter. The same session ID is given to all resources and processes that are part of a session, including Common Base Events and business processes.

To search for failed events by session ID, use the following procedure.

Procedure

1. Ensure the administrative console is running and then click **Integration Applications → Failed Event Manager**, to enter the failed event manager.
2. From the main failed event manager page, click **Search by session**.
3. Use the **Session ID** field to specify the session ID to search against.
4. Click **OK** to begin the search.

Results

The Search Results page opens, listing all failed events that originated during the specified session.

Searching for failed events by destination

Use the **By Destination** tab on the Search page to find only those failed events that are associated with a specific destination module, component, or method. The failed event manager determines the destination based on the point of failure, regardless of the type of interaction.

Before you begin

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

About this task

When performing a search, note the following:

- The values for the fields are case sensitive.
- The fields accept the asterisk (*) wildcard character.
- If you leave any field on this tab blank, the blank field is treated as a wild card. The failed event manager will search in all components, modules, or methods.
- You can search on a single destination criteria or on multiple criteria. Searching on two or more of the destination criteria provides a more refined list of failed events.

To search for failed events by destination, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Integration Applications → Failed Event Manager** to enter the failed event manager.
2. From the main failed event manager page, click **Search by destination**.

   The Search page opens with the **By Destination** tab selected.
3. Specify the search criteria you want to use. You can use any combination of the following fields to customize your search:
   - **Destination module**: use this field to specify the failed event's destination module.
   - **Destination component**: use this field to specify the failed event's destination component.
   - **Destination method**: use this field to specify the failed event's destination method.
4. Click **OK** to begin the search.

Results

The Search Results page opens and displays a list of all failed events that were destined for the specified module, component, or method.

Searching for failed events by source

Use the **By Source** tab on the Search page to find only those failed events that originated from a specific source module, component, or both. The failed event manager determines the source based on the point of failure, regardless of the type of interaction.

Before you begin

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

About this task

When performing a search, note the following:

- The values for the fields are case sensitive.
- The fields accept the asterisk (*) wildcard character.
- If you leave either field on this tab blank, the blank field is treated as a wild card. The failed event manager will search in all components or modules.
- To get the most refined list of failed events, use both the **Source module** and **Source component** fields.

To search for failed events by source, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Integration Applications → Failed Event Manager** to enter the failed event manager.
2. From the main failed event manager page, click **Search by source**.

   The Search page opens with the **By Source** tab selected.
3. Specify the search criteria. You can use one or both of the following fields:
   - **Source module**: use this field to specify the module that the failed event originated from.
   - **Source component**: use this field to specify the component that the failed event originated from.
4. Click **OK** to begin the search.

Results

The Search Results page opens and displays a list of all failed events that originated from the specified module, component, or both.

Searching for failed events by date

Use the **By Date** tab on the Search page to find only those events that failed during a specific time period.

Before you begin

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

About this task

When performing a search by date, note the following:

- The format for the date and time are locale-specific. An example of the appropriate format is provided with each field.

  Note: The values you supply must match the required format exactly. If you provide an incorrectly formatted value, 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 individual machines running the administrative console.
- You must specify a value for both fields on this tab.

To search for failed events by date, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Integration Applications → Failed Event Manager** to enter the failed event manager.
2. From the main failed event manager page, click **Search by date**.
3. Use the **From Date** field to specify the starting date and time. Because the required format for the value varies by locale, the failed event manager provides a locale-appropriate example above this field. Ensure the value you enter is formatted in the same manner as the example provided. (For instance, the required format for the en_US locale is *MM/DD/YY HH:MM Meridiem*; therefore, a correctly formatted value for this field looks like `11/10/05 4:30 PM`.)
4. Use the **To Date** field to specify the ending date and time. Because the required format for the value varies by locale, the failed event manager provides a locale-appropriate example above this field. Ensure the value you enter is formatted in the same manner as the example provided. (For instance, the required format for the en_US locale is *MM/DD/YY HH:MM Meridiem*; therefore, a correctly formatted value for this field looks like `11/17/05 4:30 PM`.)
5. Click **OK** to begin the search.

   The Search Results page opens and displays a list of all failed events that originated during the specified time period.

Searching for failed events by business object type

Use the **By Type** tab of the Search page to find only those failed events that are associated with a specific business object.

Before you begin

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

About this task

To search for failed events by business object type, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Integration Applications → Failed Event Manager** to enter the failed event manager.
2. From the main failed event manager page, click **Search by business object type**.

   The Search page opens with the **By Type** tab selected.
3. Specify the business object type you want to search against by using one of the following items:
   - **Select the business object type**: use this drop-down menu to select the type of business object associated with the failed events. This menu contains a list of all business object types found in the failed events on the server.
   - **Other business object type**: use this field to specify the type of business object associated with the failed events. The field accepts the asterisk (*) wildcard character. All values are case sensitive.
4. Click **OK** to begin the search.

Results

The Search Results page opens and displays a list of all failed events that are associated with the specified business object type.

Searching for failed events by exception

Use the **By Exception** tab of the Search page to find only those failed events that are associated with a specific exception. You can specify part or all of the exception text.

Before you begin

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

About this task

To search for failed events by exception type, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Integration Applications → Failed Event Manager** to enter the failed event manager.
2. From the main failed event manager page, click **Search by exception text**.
   The Search page opens with the **By exception** tab selected.
3. In the **Exception text** field, type the text associated with the exception you want to search against.
   You can specify all or part of the exception text, as well as the asterisk (*) wildcard character to make the search easier. The values in this field are case sensitive.

   Note: If you leave the **Exception text** field blank, it is treated as a wild card; all failed events are returned.
4. Click **OK** to begin the search.

Results

The Search Results page opens and displays a list of all failed events that are associated with the specified exception text.

Performing an advanced search for failed events

Use the **Advanced** tab of the Search page to perform a more refined search for failed events by using a combination of the criteria found on the other search tabs.

Before you begin

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

About this task

Note the following:

- Unless otherwise noted below, all fields accept the asterisk (*) wildcard character.
- Leaving a field blank causes it to be treated as a wild card.
- The advanced search is not optimized; executing an advanced search on a large set of failed events can reduce performance.

To perform an advanced search, use the following procedure.

Procedure

1. Ensure the administrative console is running, and then click **Integration Applications → Failed Event Manager** to enter the failed event manager.
2. From the main failed event manager page, click **Advanced search**.

   The Search page opens with the **Advanced** tab selected.
3. Specify the search criteria you want to use. You can use any combination of the following fields to customize your search:
   - **Destination module**: use this field to specify the failed event's destination module.
   - **Destination component**: use this field to specify the failed event's destination component.
   - **Session ID**: use this field to specify the session in which the event executed. This field does not accept the asterisk (*) wildcard character.
   - **Destination method**: use this field to specify the failed event's destination method.
   - **Source module**: use this field to specify the module that the failed event originated from.
   - **Source component**: use this field to specify the component that the failed event originated from.
   - **From Date**: use this field to specify the starting date and time if you want to search within a specific time period. This field does not accept the asterisk (*) wildcard character.
   - **To Date**: use this field to specify the ending date and time if you want to search within a specific time period. This field does not accept the asterisk (*) wildcard character.
   - **Business object type**: use this field to specify the type of business object associated with the failed events.
   - **Exception text**: use this field to specify the text associated with the exception you want to search against.
4. Click **OK** to begin the search.

Results

The Search Results page opens and displays a list of all failed events that meet the specified criteria.

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

Each failed event has the following data associated with it:

- The unique message ID and session ID for the event
- The service invocation type between SCA components
- The names of the module and component from which the event originated (the source). The failed event manager determines the source of an event based on the location where the invocation failed.
- The names of the destination module, component and method for the event. The failed event manager determines the event's destination based on the location where the invocation failed.

- The time the event failed
- The exception thrown when the event failed

This data cannot be edited. In addition, failed events can have associated trace and expiration data, both of which can be edited.

Business data

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 is 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

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

About this task

Each failed event has two types of data associated with it:

- Failed event data: information about the failed event itself, including the source and destination for the event, the time it failed, the exception it failed with, its message and session IDs, and its trace and expiration settings.
- Business data: information contained in the event. The business data can be encapsulated in a business object, or it can be simple data that is not part of a business object.

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 Message 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 you want to browse the business data associated with the failed event, click **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 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 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 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 Message 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 Process Server 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 68
   Failed events are stored in the failed event database and are retrieved through the search functionality of the failed event manager. You can search for all failed events on the all the servers within the cell, or for a specific subset of events.

Editing business data in a failed 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

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.

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 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 Message ID column) of the failed event whose data you want to edit.

   The Failed Event Details page opens.
3. From the Failed Event Details page, click **Edit business data** to access the Business Data Editor collection page.

   This page displays a hierarchical view of all of the data associated with the failed event.
4. Navigate through the business data hierarchy by clicking on the name of each parameter (these appear as links in the Parameter Name column). When you have located the parameter whose value 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 of 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 the 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 try to send an event again, you must resubmit it from the failed event manager. You can resubmit an event without changes, or 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 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 message ID to track its success or failure. If a resubmitted event fails again, it is returned to the failed event manager with its original message 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

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 event with trace

You can monitor the resubmission of a failed 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.

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.

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 68 for instructions on how to search for failed events.
2. 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.

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 68 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.

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

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, 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 on server**.

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 with reduced performance during an advanced search	"Advanced search feature is not optimized" on page 47
I am having trouble entering values in the Search page's **By Date** tab	"Values in the By Date tab automatically change to default if entered incorrectly" on page 47
I am having trouble deleting expired events	"Using the Delete Expired Events function appears to suspend the failed event manager" on page 48
I am having trouble with failed events not being created	"Failed events are not being created" on page 48

Advanced search feature is not optimized

The failed event manager's advanced search feature is not optimized. Therefore, you may experience reduced performance when using the Advanced search tab with a large set of failed events.

Values in the By Date tab automatically change to default if entered incorrectly

The Search page's **By Date** tab contains two fields: **From Date** and **To Date**. Both fields are required. The values are locale-dependendent, and they must be formatted exactly as shown in the example above the field. 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 machine 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 actually 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.

Chapter 11. 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 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.
2. 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
3. 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.

   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.
4. 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.

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

Chapter 12, "IBM Support Assistant," on page 85
The IBM Support Assistant is a tool that helps you use various IBM Support resources.

IBM Software Support Toolbar

Chapter 12. 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 V6.1* disk that is included with WebSphere Process Server and can be installed from the WebSphere Process Server launchpad. For more information, see the ″Options on the launchpad″ topic link in the Related Topics section at the bottom of this page.

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

Chapter 13, “Getting fixes,” on page 87
A product fix might be available to resolve your problem.

Chapter 11, “Searching knowledge bases,” on page 83
You can often find solutions to problems by searching IBM knowledge bases. Optimize your results by using available resources, support tools, and search methods.

Chapter 14, “Contacting IBM Software Support,” on page 89
IBM Software Support provides assistance with product defects.

Related reference

http://publib.boulder.ibm.com/infocenter/dmndhelp/v6r1mx/topic/com.ibm.websphere.wps.612.doc/doc/cins_launchpad.html

The launchpad for WebSphere Process Server provides several options you can select to install the entire server environment. This environment can include WebSphere Process Server or the WebSphere Process Server Client, WebSphere Application Server Network Deployment, a set of Web development tools, a Web server, message service clients, and additional supporting software and documentation.

Chapter 14, “Contacting IBM Software Support,” on page 89
IBM Software Support provides assistance with product defects.

 IBM Support Assistant

Chapter 13. 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. Obtain the tools required to get the fix. See ″Installing the Update Installer for WebSphere Software″.
2. 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.
3. 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.
4. 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.
5. Optional: To receive weekly notification of fixes and updates, subscribe to My Support e-mail updates.

Related reference

Chapter 12, “IBM Support Assistant,” on page 85
The IBM Support Assistant is a tool that helps you use various IBM Support resources.

Installing fix packs and refresh packs with the Update Installer

You can use the IBM Update Installer for WebSphere Software to install interim fixes, fix packs, and refresh packs collectively known as maintenance packages.

The Update Installer for WebSphere Software is also known as the update installer program, the UpdateInstaller program, and the Update Installation Wizard.

Installing the Update Installer for WebSphere Software

From the WebSphere Process Server launchpad, you can install the Update Installer for WebSphere Software, which is used to install interim fixes, fix packs and refresh packs for WebSphere Process Server.

Subscribe to My Support e-mail updates

Recommended Fixes for WebSphere Process Server

Chapter 14. 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 maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have. For information about the types of maintenance 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.

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 and a fix is delivered. IBM publishes resolved APARs on the Software Support Web site daily, so that other users who experience the same problem can benefit from the same resolution.

Related tasks

Chapter 12, “IBM Support Assistant,” on page 85
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

Chapter 12, “IBM Support Assistant,” on page 85
The IBM Support Assistant is a tool that helps you use various IBM Support resources.

IBM Software Support site

Notices

This information was developed for products and services offered in the U.S.A.

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

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

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

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

IBM World Trade Asia Corporation Licensing
2-31 Roppongi 3-chome, Minato-ku
Tokyo 106-0032, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

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

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

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

IBM Corporation
577 Airport Blvd., Suite 800
Burlingame, CA 94010
U.S.A.

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

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

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

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

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

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: (c) (your company name) (year). Portions of

this code are derived from IBM Corp. Sample Programs. (c) Copyright IBM Corp. _enter the year or years_. All rights reserved.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Programming interface information

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

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

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

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

Trademarks and service marks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation in the United States, other countries, or both. If these and other IBM trademarked terms are marked on their first occurrence in this information with a trademark symbol (® or), these symbols indicate U.S. registered or common law trademarks owned by IBM at the time this information was published. Such trademarks may also be registered or common law trademarks in other countries. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe is a registered trademark of Adobe Systems Incorporated in the United States, and/or other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

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

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

This product includes software developed by the Eclipse Project (http://www.eclipse.org).

IBM WebSphere Process Server for Multiplatforms, Version 6.1.2

Printed in USA