WebSphere® Process Server for z/OS

Version 6.0.2

IBM

Troubleshooting and Support

30 March 2007

This edition applies to version 6, release 0, modification 2 of WebSphere Process Server for z/OS (product number 5655-N53) and to all subsequent releases and modifications until otherwise indicated in new editions.

To send us your comments about this document, email 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. Troubleshooting overview

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, even though one layer might report the problem, this does not mean that the problem originates 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, its version, all corresponding software and versions, and hardware information. Confirm that you are running within an environment that is a supported

1

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, try to answer these 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 questions like this can help to 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, 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.

These and other 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. Debugging applications

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

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 IBM® WebSphere® Integration Developer.

Chapter 3. Adding logging and tracing to your applications

Designers and developers of applications that run on WebSphere Process Server can find it useful to use Java logging for generating their application logging.

IBM WebSphere Process Server is built on IBM WebSphere Application Server, Network Deployment and also works with infrastructure and platform services from IBM WebSphere Application Server. For more information, see the **Logging and tracing with Java logging** and **The Common Base Event in WebSphere Application Server** topics in the WebSphere Application Server for z/OS information center.

1. Follow the instructions in Logging and tracing with Java logging.
2. Follow the instructions in The Common Base Event in WebSphere Application Server.
3. In addition, WebSphere Process Server monitoring capabilities use logging. For more information, see the WebSphere Process Server *Monitoring* PDF file.

Chapter 4. Managing WebSphere Process Server failed events

TheWebSphere 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.

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**.

Related concepts

"Working with data in failed events" on page 16
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.

"Resubmitting failed events" on page 19
If you want to try to execute the 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.

Related tasks

"Finding failed events" on page 11
Failed events are stored in the failed event database and are retrieved through the failed event manager's search functionality. You can search for all failed events on the all the servers within the cell, or for a specific subset of events.

"Deleting failed events" on page 21
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.

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 the *Securing applications and their environment* topic in the WebSphere Process Server Information Center.

Finding failed events

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

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

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

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

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

Related tasks

"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.

"Searching for failed events by source" on page 13
Use the Search page's **By Source** tab 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.

"Searching for failed events by destination" on page 12
Use the Search page's **By Destination** tab 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.

"Searching for failed events by date" on page 13
Use the Search page's **By Date** tab to find only those events that failed during a specific time period.

"Searching for failed events by business object type" on page 14
Use the Search page's **By Type** tab to find only those failed events that are associated with a specific business object.

"Searching for failed events by exception" on page 14
Use the Search page's **By Exception** tab to find only those failed events that are associated with a specific exception. You can specify part or all of the exception text.

"Performing an advanced search for failed events" on page 15
Use the Search page's **Advanced** tab to perform a more refined search for failed events by using a combination of the criteria found on the other search tabs.

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.

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

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.

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

Searching for failed events by destination

Use the Search page's **By Destination** tab 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.

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

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:
   - The **Destination module** field—Use this field to specify the failed event's destination module.
   - The **Destination component** field—Use this field to specify the failed event's destination component.
   - The **Destination method** field—Use this field to specify the failed event's destination method.
4. Click **OK** to begin the search.

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 Search page's **By Source** tab 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.

You must be logged in as administrator or operator to perform 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.
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:
   - The **Source module** field—Use this field to specify the module that the failed event originated from.
   - The **Source component** field—Use this field to specify the component that the failed event originated from.
4. Click **OK** to begin the search.

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 Search page's **By Date** tab to find only those events that failed during a specific time period.

You must be logged in as administrator or operator to perform 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.
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 Search page's **By Type** tab to find only those failed events that are associated with a specific business object.

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

To search for failed events by business object type, use the following 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:
   - The **Select the business object type** menu—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.
   - The **Other business object type** field—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.

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

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

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

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

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

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:
   - The **Destination module** field—Use this field to specify the failed event's destination module.
   - The **Destination component** field—Use this field to specify the failed event's destination component.
   - The **Session ID** field – Use this field to specify the session in which the event executed. This field does not accept the asterisk (*) wildcard character.
   - The **Destination method** field—Use this field to specify the failed event's destination method.
   - The **Source module** field—Use this field to specify the module that the failed event originated from.
   - The **Source component** field—Use this field to specify the component that the failed event originated from.

- The **From Date** field—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.
- The **To Date** field—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.
- The **Business object type** field—Use this field to specify the type of business object associated with the failed events.
- The **Exception text** field—Use this field to specify the text associated with the exception you want to search against.
4. Click **OK** to begin the search.

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.

You must be logged as administrator or operator to perform 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.

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 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 paramter is a complex data type, clicking its name will expand the hierachy 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.

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

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 11
   Failed events are stored in the failed event database and are retrieved through the failed event manager's search functionality. 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.

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

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 execute the 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.

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

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**.

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 executes successfully. The failed event manager provides optional tracing for all failed events.

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 WebSphere Process Server Information Center.

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

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.

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.

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

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 11 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**.

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.

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

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 11 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**.

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.

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

To delete one or more failed events, use the following 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"
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	"Executing the Delete Expired Events function appears to suspend the failed event manager"
I am having trouble with failed events not being created	"Failed events are not being created" on page 23

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.

Executing 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 5. Diagnosing problems

You can use problem determination to understand why your application, or server, is not working.

Diagnosing problems can aid you in understanding why your enterprise application or server is not working and can help you resolve problems. Unlike performance tuning, which focuses on solving problems associated with slow processes and un-optimized performance, problem determination focuses on finding solutions to functional problems. For more information about diagnosing problems, see the **Diagnosing problems (using diagnosis tools)** section in the WebSphere Application Server Network Deployment documentation.

1. Investigate common problems organized according to functional areas within IBM WebSphere Process Server in Troubleshooting by function.

2. If you already have an error message and want to quickly look up its explanation and recommended response, see the Messages section of the WebSphere Process Server documentation.

3. For help in knowing where to find error and warning messages, interpreting messages, and configuring log files, expand the **Diagnosing problems (using diagnosis tools)** section in the WebSphere Application Server Network Deployment documentation table of contents and select **Diagnosing problems with message logs**.

4. 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, expand the **Diagnosing problems (using diagnosis tools)** section in the WebSphere Application Server Network Deployment documentation table of contents and select **Working with trace**.

5. For help in adding log and trace capability to your own application, see Adding logging and tracing to your application.

6. For help in using settings or tools to help you diagnose the problem, expand **Diagnosing problems (using diagnosis tools)** section in the WebSphere Application Server Network Deployment documentation table of contents and select **Working with troubleshooting tools**. Some of these tools are bundled with the product, and others are downloadable.

7. To find out how to look up documented problems, common mistakes, WebSphere Process Server prerequisites, and other problem-determination information available on the WebSphere Process Server public Web site, or to obtain technical support from IBM, see Obtaining help from IBM.

8. The IBM Developer Kit and Runtime Environment, Java™ 2 Technology Edition, Version 1.4.2 Diagnostics Guide describes debugging techniques and the diagnostic tools that are available to help you solve problems with Java. It also gives guidance on how to submit problems to IBM. You can find the guide at IBM developer kits: Diagnosis documentation Web site.

9. For current information available from IBM Support on known problems with WebSphere Process Server and their resolutions, see the WebSphere Process Server support page.

Chapter 6. Troubleshooting by function

Depending on the specific problem that you have encountered, it might be necessary to troubleshoot by function. You can troubleshoot based on different issues that are common in installation, configuration, deployment, and administration.

Read the Chapter 1, "Troubleshooting overview," on page 1 before you begin troubleshooting by function.

Troubleshooting the installation and configuration

If the product installation and configuration are not successful, use this information to help you assess and correct the problems.

This topic assumes that you have attempted to install and or configure but have not been successful.

You should be aware that the installer program records multiple indicators of success in the following ways:

- Standard output messages

  Standard output messages display directly on the screen from which you run the installation script (zSMPInstall.sh), or configuration script(s) zWPSConfig.sh / zWESBConfig.sh. You can choose to *redirect* these messages to a file by using redirect symbol and a file name at the end of the command line. For example, adding the syntax >run.log to the end of the installation command will redirect the standard output messages to the file named **run.log** in the present working directory.

- Log file messages

  Log messages for installation are written to the **zSMPInstall.log** file in the runtime directory. Standard location for this file is /WebSphere/V6R0M0/ AppServer/logs/wbi/zSMPInstall.log.

  Log messages for configuration are written to the **zWPSConfig.log** or the **zWESBConfig.log** file in the runtime directory. Standard locations for these files are /WebSphere/V6R0M0/AppServer/logs/wbi/zWPSConfig.log and /WebSphere/V6R0M0/AppServer/logs/wbi/zWESBConfig.log respectively.

- Trace file messages

  The installation trace messages are written to the **zSMPInstall.trace** file in the runtime directory. Standard location for this file is /WebSphere/V6R0M0/ AppServer/logs/wbi/zSMPInstall.trace.

  Trace messages for configuration are written to the **zWPSConfig.trace** or the **zWESBConfig.trace** file in the runtime directory. Standard locations for these files are /WebSphere/V6R0M0/AppServer/logs/wbi/zWPSConfig.trace and /WebSphere/V6R0M0/AppServer/logs/wbi/zWESBConfig.trace respectively.

Make sure that you have installed and configured WebSphere Application Server for z/OS® successfully. Refer to the installation troubleshooting information in the WebSphere Application Server for z/OS information center if you are having trouble installing and configuring WebSphere Application Server for z/OS.

For current information available from IBM Support on known problems and their resolution, see the WebSphere Process Server Support page.

1. Review the messages from Standard Out. There should be no error messages displayed. The standard output messages display on either the screen from which you ran the installation command or in a file that you specified by using the redirect (">") symbol on the command line.

   The following is an example of a successful execution of the installation script with the **-install** option:

   ```
   parsing command arguments...
   parsing arguments complete
   setting up configuration...
   runtimeRootDirName is: /WebSphere/V6R0M0/AppServer
   WAS_HOME is: /WebSphere/V6R0M0/AppServer
   WBI_HOME is: /WebSphere/V6R0M0/AppServer
   set up configuration complete
   creating the symbolic links...
   invokeSymLink
   creation of symbolic links complete
   doing post install file updates...
   post install updates complete
   running Configuration Manager update...
   Configuration Manager update complete
   augmenting profile(s)...
   augmenting profile(s) complete
   ```

   Error messages indicate an unsuccessful installation. Some errors as displayed in Standard Out will be self explanatory and can be easily corrected. If the installation completed to the point where the log and trace files were created, continue with the following steps.

2. Review the zSMPInstall.log (ASCII) file in the runtime directory. For a stand-alone configuration, the standard location for this file is /WebSphere/V6R0M0/AppServer/logs/wbi/zSMPInstall.log. For a deployment manager configuration, the standard location for this file is /WebSphere/V6R0M0/DeploymentManager/logs/wbi/zSMPInstall.log

   If there are error messages, try to determine which of the following tasks were in progress when the error occurred.

   - create symbolic links
   - create post installation file
   - update codebase permissions
   - update Configuration Manager

   Knowing the task that was in progress at the time of an error will help you assess the information in the trace file.

3. Review the zSMPInstall.trace or the zWPSConfig.trace / zWESBConfig.trace (ASCII) file in the runtime directory. For a stand-alone configuration, the standard location for these files are /WebSphere/V6R0M0/AppServer/logs/wbi/zSMPInstall.trace or /WebSphere/V6R0M0/AppServer/logs/wbi/zWPSConfig.trace For a deployment manager configuration, the standard location of this file is /WebSphere/V6R0M0/DeploymentManger/logs/wbi/zSMPInstall.trace or /WebSphere/V6R0M0/DeploymentManger/logs/wbi/zWPSConfig.trace

   On a successful installation, only informational messages (messages with a suffix of **I**, for example CWPIZ0044**I**) should be listed in the trace file.

   If warning messages (messages with a suffix of **W**) or error messages (messages with a suffix of **E**) are listed in the trace, further review is required.

If the warning or error occurred <u>during the create symbolic links, create post installation file, or update codebase permissions tasks</u>, the trace message should contain information that will help you diagnose and correct the problem.

If the warning or error occurred in the update Configuration Manager task, proceed to the next step.

If the warning or error occurred in the augment profile(s) task, proceed to step 5.

4. Review the actions of the Update Configuration Manager task. These actions are recorded by writing to a log file (ASCII). The log file name is `cmtInstall.log`.

Standard location for this file is in directory /WebSphere/V6R0M0/AppServer/logs/wbi.

Search the Configuration Manager log for `>SEVERE<` or `>WARNING<` level messages to determine overall error in processing.

Each Ant script run from the installation directory writes to its own log (ASCII).

To determine what Ant file was running at the time of the error, you can look for 'Buildfile' previous.

The default name for the installation directory that contains the ant scripts is: /WebSphere/V6R0M0/AppServer/properties/version/install.wbi/6.0.0.0/config/full/install.

The resulting ant logs are written to the product log directory. The default name for this directory is /WebSphere/V6R0M0/AppServer/logs/wbi. Ant logs include the following (review these logs to determine errors in processing) :

```
100SUgradeCoreAdminConsolePlugins.ant.log
101SUgradeServerAdminConsolePlugins.ant.log
101SWbiWebuiUpgrade.ant.log
102SUpgradeIscdeploy.ant.log
90SConfigNoProfileFirstStepsESB.ant.log
90SConfigNoProfileFirstStepsWBI.ant.log
90SConfigureWSProfileForWBI.ant.log
90SInstallCEI.ant.log
90SUpdateJavaOptions.ant.log
91SConfigNoProfileFirstStepsCharset.ant.log
98SDeployBPCAdminConsolePlugins.ant.log
98SDeployServerAdminConsolePlugins.ant.log
99SDeployCoreAdminConsolePlugins.ant.log
99WbiProfileUpgrade.log
```

If there were no problems a **BUILD SUCCESSFUL** message displays at the end of the file.

5. Review the actions of the WebSphere Application Server profile augment task. The augment profile(s) task records its actions by writing to a log file (ASCII). The log file name has the name **wasprofile_augment_default.log**. Standard location for this file is in the directory /WebSphere/V6R0M0/AppServer/logs/wasprofile.

Search the Profile Augmentation log for `>SEVERE<` or `>WARNING<` level messages to determine overall error in processing.

Individual Ant action logs are located in /WebSphere/V6R0M0/AppServer/profiles/default/logs.

To determine what Ant file was running at the time of the error, you can look for 'Buildfile' previous.

After troubleshooting the problems that caused the installation errors, and after you have run the installation script successfully, you should perform the following steps:

1. Start the WebSphere Application Server.
2. Launch the Administrative Console and verify that the product components have been installed.

   For example, by installing WebSphere Process Server successfully, you should see evidence of Process Choreographer under the Enterprise Applications with names that start with `BPEContainer`, `BPCExplorer`, and `TaskContainer`.

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

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

About the installation error messages

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

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

The WebSphere Process Server for z/OS installation error messages are documented in the information center under **Reference > Messages >CWPIZ**.

The WebSphere Process Server for z/OS installation error messages are written to the zSMPInstall.log file in the run-time directory. The standard default location for the log file is /WebSphere/V6R0M0/AppServer/logs/wbi/zSMPInstall.log.

The WebSphere Process Server for z/OS configuration error messages are written to the zWPSConfig.log file and the zWESBConfig.log file in the run-time directory. The standard default location for these log files are /WebSphere/V6R0M0/ AppServer/logs/wbi/zWESBConfig.log and /WebSphere/V6R0M0/AppServer/ logs/wbi/zWPSConfig.log respectively.

Log files

Various log files are created during the product installation and configuration process.

Purpose

Consult the applicable logs if problems occur during the product installation and configuration process.

Standard out messages redirected to log file

Standard out messages report high-level actions such as the starting and completing of the action that verifies the command line arguments.

By default, these messages display directly on the screen from which you run the product installation script. However, you can *redirect* these messages to a file by using the redirect symbol and a file name at the end of the command line. For example, specifying >run.log at the end of the installation command redirects the standard out messages to a file named run.log in the present working directory.

Standard out messages also report severe errors that occur prior to the Log and Trace File being opened. For instance, the following message block displays if a required keyword (-runtime) was not included in the installation command.

```
parsing command arguments...
CWPIZ0101E -runtime keyword and value not specified on command line.
com.ibm.ws390.installer.InstallFailureException: -runtime keyword and value not specified
CWPIZ0017E  install task failed.
```

Log file

These messages include the messages written to Standard Out, but provide additional information and settings that were used by the installer program.

For instance, the following log portion shows the response properties and their values being used. It also shows the source and target directories being used during the creation of the symbolic links.

```
response property: profilePath=/WebSphere/V6R0M0/AppServer/profiles/default
response property: nodeName=SY1
response property: scaSecurityPassword=ibmuser
response property: dbType=Cloudscape
response property: ceiSampleJmsUser=ibmuser
response property: scaSecurityUserId=ibmuser
response property: configureScaSecurity=true
response property: mqUser=ibmuser
response property: serverName=server1
response property: adminBFMGroups=ibmuser
response property: profileName=default
response property: dbCreateNew=true
response property: ceiSampleJmsPwd=ibmuser
response property: cellName=SY1
response property: dbLocation=/WebSphere/V6R0M0/AppServer/cloudscape/databases/WBIDB
response property: mqPwd=ibmuser
response property: was.install.root=/WebSphere/V6R0M0/AppServer
response property: augment=
response property: ceiDbProduct=CLOUDSCAPE_V51_1
response property: wbi.install.root=/WebSphere/V6R0M0/AppServer
response property: ceiSampleServerName=server1
response property: templatePath=/WebSphere/V6R0M0/AppServer/profileTemplates/default.*
response property: dbName=WBIDB
set up configuration complete
creating the symbolic links...
Source=/usr/lpp/zWPS/V6R0M0

Target=/WebSphere/V6R0M0/AppServer
creation of symbolic links complete
doing post install file updates...
post install updates complete
running Configuration Manager update...
Configuration Manager update complete
```

Trace file

These messages are written to the **zSMPInstall.trace** file in the run-time directory.

The example below shows some preliminary informational messages and then a
CWPIZ0322E error indicating that the required profileName property was not
found in the response file that the user specified on the installation script
command line (nor was provided as a -Z override).

The subsequent **CWPIZ0017E** error message is a general message indicating the
final outcome of the zSMPInstall.sh run.

```
[8/16/05 17:00:45:380 EDT] 0000000a ManagerAdmin  I   BB000222I:
TRAS0017I: The startup trace state is *=info.

[8/16/05 17:00:48:230 EDT] 0000000a WPSInstaller  I   BB000222I:
CWPIZ0044I Begin install task.

[8/16/05 17:00:48:273 EDT] 0000000a WPSInstaller  I   BB000222I:
CWPIZ0117I WPS installer log data will be written to
/WebSphere/V6R0M0/AppServer/logs/wbi/zSMPInstall.log.

[8/16/05 17:00:48:282 EDT] 0000000a WPSInstaller  I   BB000222I:
CWPIZ0024I WPS installer trace data will be written to
/WebSphere/V6R0M0/AppServer/logs/wbi/zSMPInstall.trace.

[8/16/05 17:00:48:292 EDT] 0000000a WPSInstaller  I   BB000222I:
CWPIZ0014I Trace specification is "*=all=disabled".

[8/16/05 17:00:48:298 EDT] 0000000a WPSInstaller  I   BB000222I:
CWPIZ0045I WPS SMP/E root directory is /zrockuser/wbi/Install.

[8/16/05 17:00:48:302 EDT] 0000000a WPSInstaller  I   BB000222I:
CWPIZ0052I WAS SMP/E root directory is /web/usr/lpp/zWebSphere/V6R0.

[8/16/05 17:00:48:307 EDT] 0000000a WPSInstaller  I   BB000222I:
CWPIZ0046I Destination application server root directory is
/WebSphere/V6R0M0/AppServer.

[8/16/05 17:00:48:314 EDT] 0000000a WPSInstaller  E   BB000220E:
CWPIZ0322E profileName property not specified in Response File.

[8/16/05 17:00:48:318 EDT] 0000000a WPSInstaller  E   BB000220E:
CWPIZ0017E install task failed.

A trace file from a zSMPInstall.sh executed with the trace specification argument
set to "*=all=enabled" provides additional debugging information. It may contain
information that is meaningful only to a developer. The following is a partial
trace using "*=all=enabled":


************ Start Display Current Environment ************
Host Operating System is z/OS, version 01.04.00
Java version = J2RE 1.4.2 IBM z/OS Persistent Reusable VM build cm142-20050623
(JIT enabled: jitc), Java
Compiler = jitc, Java VM name = Classic VM
was.install.root = /WebSphere/V6R0M0/AppServer
user.install.root = /WebSphere/V6R0M0/AppServer/profiles/default
Java Home = /web/usr/lpp/zWebSphere/V6R0/java/J1.4
ws.ext.dirs = /WebSphere/V6R0M0/AppServer/java/lib:/WebSphere/V6R0M0/AppServer/java/lib/
ext:/WebSphere/V6R0M0/AppServer/classes:/WebSphere/V6R0M0/AppServer/lib:/WebSphere/V6R0M0/AppServer/
installedChannels:/WebSphere/V6R0M0/AppServer/lib/ext:/WebSphere/V6R0M0/AppServer/deploytool/itp
/plugins/com.ibm.etools.ejbdeploy/runtime:/WebSphere/V6R0M0/AppServer/MQSeries/pubsubroot/lib
Classpath = /zrockuser/bbzconfig.jar:/WebSphere/V6R0M0/AppServer/admin.jar:/WebSphere/V6R0M0
/AppServer/lib/ant.jar:/WebSphere/V6R0M0/AppServer/lib/bootstrapws390.jar:/WebSphere/V6R0M0
/AppServer/lib/bootstrap.jar:/WebSphere/V6R0M0/AppServer/lib/configmanager.jar:/WebSphere
/V6R0M0/AppServer/lib/emf.jar:/WebSphere/V6R0M0/AppServer/lib/ras.jar:/WebSphere/V6R0M0
/AppServer/lib/runtimefw.jar:/WebSphere/V6R0M0/AppServer/lib/utils.jar:/WebSphere/V6R0M0
/AppServer/lib/wasjmx.jar:/WebSphere/V6R0M0/AppServer/lib/wasproduct.jar:/WebSphere/V6R0M0
/AppServer/lib/wccm_base.jar:/WebSphere/V6R0M0/AppServer/lib/wjmxapp.jar:/WebSphere/V6R0M0
/AppServer/lib/wsanttasks.jar:/WebSphere/V6R0M0/AppServer/lib/wsexception.jar:/WebSphere
/V6R0M0/AppServer/lib/wsprofile.jar:/WebSphere/V6R0M0/AppServer/profiles/default/properties:
/WebSphere/V6R0M0/AppServer/properties:/WebSphere/V6R0M0/AppServer/lib/bootstrap.jar:/WebSphere
/V6R0M0/AppServer/lib/j2ee.jar:/WebSphere/V6R0M0/AppServer/lib/lmproxy.jar:/WebSphere/V6R0M0
/AppServer/lib/urlprotocols.jar:/WebSphere/V6R0M0/AppServer/lib/bootstrapws390.jar
Java Library path = /web/usr/lpp/zWebSphere/V6R0/java/J1.4/bin/classic/libjvm.so:/web/usr
/lpp/zWebSphere/V6R0/java/J1.4/bin/classic:/web/usr/lpp/zWebSphere/V6R0/java/J1.4/bin/:
/WebSphere/V6R0M0/AppServer/lib:/WebSphere/V6R0M0/AppServer/lib:/WebSphere/V6R0M0/AppServer
/MQSeries/pubsubroot/lib:/mqm/java/bin:/mqm/java/lib:/db2810/lib:/db2beta/db2710/lib:
/web/usr/lpp/WebSphere/lib:/lib:/usr/lib:/java/J1.3/bin:/java/J1.4/bin:/java/J5.0/bin:
/staf/lib:/WebSphere/V6R0M0/AppServer/lib:/usr/lib
Current trace specification = *=all
************ End Display Current Environment ************
[10/3/05 16:35:05:709 EDT] 0000000a ManagerAdmin  I   BB000222I: TRAS0017I:
The startup trace state is *=all.
[10/3/05 16:35:08:638 EDT] 0000000a WPSInstaller  >  setup Entry
                             /web/usr/wbi/zWebSphere/V6R0
                             APPSERVER
                             zSMPInstall.sh
                             -smproot
                             /web/usr/wbi/zWPS/V6R0
```

```
                     -runtime
                     /WebSphere/V6R0M0/AppServer
                     -response
                     /web/usr/wbi/zWPS/V6R0/zos.config/standAloneProfile.rsp
                     -prereqonly
                     -trace
                     *=all=enabled
[10/3/05 16:35:08:640 EDT] 0000000a WPSInstaller  3   logFileDeleted
                     true
[10/3/05 16:35:08:660 EDT] 0000000a WPSInstaller  I   BBOO0222I: CWPIZ0044I:
Begin install task.
[10/3/05 16:35:08:702 EDT] 0000000a WPSInstaller  I   BBOO0222I: CWPIZ0117I:
WPS installer log data will be written to /WebSphere/V6R0M0/AppServer/logs/wbi/zSMPInstall.log.
[10/3/05 16:35:08:712 EDT] 0000000a WPSInstaller  I   BBOO0222I: CWPIZ0024I:
WPS installer trace data will be written to /WebSphere/V6R0M0/AppServer/logs/wbi/zSMPInstall.trace.
[10/3/05 16:35:08:722 EDT] 0000000a WPSInstaller  I   BBOO0222I: CWPIZ0014I:
Trace specification is "*=all=enabled".
[10/3/05 16:35:08:726 EDT] 0000000a WPSInstaller  I   BBOO0222I: CWPIZ0052I:
WAS SMP/E root directory is /web/usr/lpp/zWebSphere/V6R0.
[10/3/05 16:35:08:730 EDT] 0000000a WPSInstaller  >   checkPathName Entry
                     /web/usr/wbi/zWPS/V6R0
[10/3/05 16:35:08:731 EDT] 0000000a WPSInstaller  <   checkPathName Exit
[10/3/05 16:35:08:732 EDT] 0000000a WPSInstaller  I   BBOO0222I: CWPIZ0045I:
WPS SMP/E root directory is /web/usr/wbi/zWPS/V6R0.
[10/3/05 16:35:08:736 EDT] 0000000a Symlink       >   isSymlink Entry
                     /web/usr/wbi/zWPS/V6R0
[10/3/05 16:35:08:737 EDT] 0000000a Symlink       3   absolute path
                     /web/usr/wbi/zWPS/V6R0
[10/3/05 16:35:08:737 EDT] 0000000a Symlink       3   canonical path
                     /web/usr/wbi/zWPS/V6R0
[10/3/05 16:35:08:738 EDT] 0000000a Symlink       <   isSymlink Exit
                     false
[10/3/05 16:35:08:738 EDT] 0000000a WPSInstaller  I   BBOO0222I: CWPIZ0046I:
Destination application server root directory is /WebSphere/V6R0M0/AppServer.
[10/3/05 16:35:08:744 EDT] 0000000a WPSInstaller  I   BBOO0222I: CWPIZ0247I:
Response file is /web/usr/wbi/zWPS/V6R0/zos.config/sample.rsp.
[10/3/05 16:35:08:764 EDT] 0000000a WPSInstaller  3   response property
                     profilePath=/WebSphere/V6R0M0/AppServer/profiles/default
[10/3/05 16:35:08:765 EDT] 0000000a WPSInstaller  3   response property
                     nodeName=SY1
```

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.

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.

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.

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 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. Refer to **Developing and deploying modules > Overview of preparing and installing modules > Preparing to deploy to a server** for 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.

   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.

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.

Delete J2C activation specifications when you inadvertently saved a configuration after installing an application that contains services and do not require the specifications.

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 stand alone server or at the **server** scope in a Network 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.**

The system removes selected specifications from the display.

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.

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.

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.

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.**

The panel displays only the remaining destinations.

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

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.

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 22
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 22
I am having trouble deleting expired events	"Executing the Delete Expired Events function appears to suspend the failed event manager" on page 22
I am having trouble with failed events not being created	"Failed events are not being created" on page 23

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.

Executing 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.

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.

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.

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 **Resources > JMS Providers > WebSphere MQ > WebSphere MQ queue destinations**.
4. Locate any messages that may indicate failed transactions. You will have to investigate system, SCA-specific message areas, 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 MQ error, there will be an MQException linked somewhere in the stack trace with an 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) applications.
   d. Examine the 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 "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.

WebSphere Application Server troubleshooting

Because IBM WebSphere Process Server is built on IBM WebSphere Application Server, you may want to consult troubleshooting information in the WebSphere Application Server documentation.

IBM WebSphere Process Server is built on WebSphere Application Server Network Deployment, version 6.0.2. WebSphere Process Server also works with infrastructure and platform services from IBM WebSphere Application Server, version 6.0.2.

For more information about troubleshooting in WebSphere Application Server, see Troubleshooting and support in the WebSphere Application Server Network Deployment documentation. To view the topic, expand **Troubleshooting and support > Troubleshooting WebSphere applications** in the table of contents.

Chapter 7. Resources for diagnosing and fixing problems

In addition to the information center, there are several Web-based resources for researching and resolving problems related to WebSphere Process Server.

Product support page: The official site for providing tools and sharing knowledge about problems with WebSphere Process Server, is the WebSphere Process Server support page.

The support page includes the following resources and capabilities:
* A search field for searching the entire support site for documentation and fixes related to a specific exception, error message, or other problem. Use this search function before contacting IBM Support directly.
* **Hints and Tips**, **Technotes**, and **Solutions** links take you to specific problems and resolutions documented by WebSphere Process Server technical support personnel.
* Links provide access to free WebSphere Process Server maintenance upgrades and problem determination tools.
  – Fixes are software patches which address specific WebSphere Process Server defects. Selecting a specific defect from the list in the **Fixes by version** page takes you to a description of what problem the fix addresses.
  – Tools are free programs that help you analyze the configuration, behavior and performance of your WebSphere Process Server installation.

Note: Some resources on the WebSphere Process Server support page are marked with a key icon. To access these resources, you must supply a user ID and password, or to register if you do not already have an ID. When registering, you are asked for your contract number, which is supplied as part of a WebSphere Process Server purchase.

WebSphere developerWorks: IBM developerWorks® is an IBM-supported site for enabling developers to learn about IBM software products and how to use them. They contain resources such as articles, tutorials, and links to newsgroups and user groups. For more information, see IBM developerWorks Web site. Additional information about WebSphere Process Server can be found on the IBM developerWorks Business Integration Zone.

IBM Support page: IBM Support provides documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the IBM Support page.

IBM Support Assistant

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

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.
* an Education component, which provides guided access to IBM product education web sites, including IBM Education Assistant modules.

- a Support Links component, which provides a convenient location to access various IBM Web resources such as IBM product sites, IBM support sites and links to IBM news groups.
- 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, version 2.0, and then installing plug-ins for WebSphere Process Server.

For more information, see Support Assistant product plug-in for WebSphere Process Server.

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

If you already have IBM Support Assistant installed, you can download a plug-in for WebSphere Process Server from within IBM Support Assistant by clicking **Find Plug-ins** and then clicking **WebSphere** to initiate a search for all WebSphere product plug-ins. When you click the link for the WebSphere Process Server plug-in, the download page opens.

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

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.

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
   - 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. For more information, see the Support Assistant product plug-in for WebSphere Process Server.

   If you cannot find an answer to your question in the information center, search the Internet for the latest, most complete information that might help you

resolve your problem. To search multiple Internet resources for your product, open the IBM Support Assistant and click **Web Search**. From this page, you can search a variety of resources including:

- IBM technotes
- IBM downloads
- IBM books and articles
- IBM developerWorks
- IBM newsgroups and forums
- Google

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

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 messages, their explanations, and the recommended recovery actions by searching for the message identifier in the Reference section of the WebSphere Process Server documentation.

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.

Obtaining help from IBM

If you are not able to resolve a problem with IBM WebSphere Process Server by following the steps described in the Troubleshooting documentation, by looking up error messages in the message reference, or by looking for related documentation on the online help, contact IBM Technical Support.

Purchase of IBM WebSphere Process Server entitles you to one year of telephone support under the Passport Advantage® program. For details on the Passport Advantage program, visit Passport Advantage Web site.

The number for Passport Advantage members to call for WebSphere Process Server support is 1-800-237-5511. Have the following information available when you call:

- Your Contract or Passport Advantage number.
- Your WebSphere Process Server version and revision level, plus any installed fixes.
- Your operating system name and version.
- Your database type and version.
- Basic topology data: how many machines are running how many application servers, and so on.
- Any error or warning messages related to your problem.

IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the WebSphere Process Server MustGather technote, which is available from the IBM Support page for WebSphere Process Server.

For more information, see Contacting IBM Software Support.

1. Enable tracing. WebSphere Process Server support engineers might ask you to enable tracing on a particular component of the product to diagnose a difficult problem.

2. Use consulting services. For complex issues such as high availability and integration with legacy systems, education, and help in getting started quickly with the WebSphere product family, consider using IBM consulting services. To learn about these services, browse the IBM Global Services Web site.

Chapter 8. Getting fixes

A product fix might be available to resolve your problem.

1. Determine which fix you need. Check the list of WebSphere Process Server recommended fixes to confirm that your software is at the latest maintenance level. Check the list of problems fixed in WebSphere Process Server fix readme documentation 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.

2. Download the fix. Open the download document and follow the link in the **Download package** section.

3. Apply the fix. Follow the instructions in the **Installation Instructions** section of the download document. For more information, see Installing maintenance packages.

4. **Optional:** To receive weekly notification of fixes and updates, subscribe to My Support e-mail updates.

Chapter 9. Applying product maintenance

Because WebSphere Process Server for z/OS is installed and configured into the WebSphere Application Server, the maintenance applied to WebSphere Process Server for z/OS is done so through the WebSphere Application Server product, using the WebSphere Application Server techniques for applying product maintenance.

Contact the IBM Software Support Center for information about preventive service planning (PSP) upgrades for the product. For more information about PSP upgrades for WebSphere Process Server for z/OS, see the *Program Directory for WebSphere Process Server for z/OS*. Although the Program Directory contains a list of required program temporary fixes (PTFs), the most current information is available from the IBM Software Support Center.

Use the following procedure whenever you want to apply a new service release to your system.

See *Applying product maintenance* in the WebSphere Application Server for z/OS information center for a description of how to apply product maintenance

You can maintain service to clients when upgrading the host cluster of WebSphere Application Server for z/OS.

Applying a service level or restoring to the previous accepted service level

Because WebSphere Process Server for z/OS is installed and configured into the WebSphere Application Server, the service level applied to WebSphere Process Server for z/OS is done so through the WebSphere Application Server product, using the WebSphere Application Server techniques for applying service level or restoring to the previous accepted service level

Service that is applied to the product data sets and product HFS occasionally requires corresponding changes to be made to the configuration HFS for existing application serving environments that configure at a lower service level. Most of these ″post-maintenance″ or ″post-install″ updates can be performed automatically. This is done by the post-installer. See *Applying a service level or restoring to the previous accepted service level* in the WebSphere Application Server for z/OS information center for a description of how to apply service

Chapter 10. Contacting IBM Software Support

IBM Software Support provides assistance with product defects.

To take advantage of unique Support features, see the WebSphere Process Server support page.

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* at techsupport.services.ibm.com/guides/services.html.

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* at techsupport.services.ibm.com/guides/beforecontacting.html.

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.

3. Submit your problem to IBM Software Support in one of the following ways:
   - Using IBM Support Assistant: See Support Assistant product plug-in for WebSphere Process Server.
   - Online: Click **Submit and track problems** on the IBM Software Support site.
   - 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.

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: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © 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 and related trademarks: http://www.ibm.com/legal/copytrade.shtml

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/).

WebSphere Process Server for z/OS, Version 6.0.2

IBM ®