



Problem Determination

Note

Before using this information, be sure to read the general information under “Notices” on page 193.

Compilation date: May 20, 2004

© Copyright International Business Machines Corporation 2003, 2004. All rights reserved.

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

Contents

How to send your comments v

Chapter 1. Troubleshooting or problem determination 1

Chapter 2. Diagnosing and fixing problems 3

Acquiring skills for problem determination 3
Working with diagnostic tools and controls 4

Chapter 3. Adding logging and tracing to your application 7

Programming with the JRas framework 7
 Understanding the JRas facility 7
 JRas Extensions 9
 JRas Messages and Trace event types 19
 Instrumenting an application with JRas extensions 22
Logging messages and trace data for Java server applications 28
 Determining where to issue the messages 28
 System performance when logging messages and trace data 28
 Issuing application messages to the MVS master console 29

Chapter 4. Best practices for maintaining the run-time environment . 31

Best practices for using system controls 32
Updating the CFRM policy 32
Configuring WebSphere Application Server for z/OS variables 33
 Steps for configuring WebSphere variables 33
Setting up component trace (CTRACE) 34
 Steps for preparing CTRACE controls and resources 35
 Steps for starting CTRACE as part of WebSphere Application Server for z/OS customization 37
 Steps for starting CTRACE while WebSphere Application Server for z/OS servers are active. 38
 Using CTRACE to collect trace data for Java server applications 38
 Setting up the error log 39
 Using the z/OS modify command. 39
Types of configuration variables 44
 Setting output destinations and characteristics. 44
 Setting trace controls 45
 Setting dump controls 46
 Controlling behavior through timeout values 47
Debugging WebSphere Application Server applications 53
 Attaching WebSphere Studio Application Developer to a remote debug session. 54
 Unit testing with DB2 55

Chapter 5. Viewing diagnostic information 57

Viewing CEEDUMPs in the job log 57
Viewing SVC dumps 57
Viewing CTRACE and JRas data through IPCS 58
 Steps for using the IPCS dialog to format CTRACE data 58
 Finding the subname for IPCS CTRACE. 59
 Steps for using IPCS in batch mode to format CTRACE data 60
Viewing error log contents through the Log Browse Utility (BBORBLOG) 62
 Using the log browse utility (BBORBLOG) 63
 Error log stream record output 64
Using the z/OS display command. 65
Converting Java minor codes 66
Viewing SYSPRINT output, STDOUT and STDERR streams in an HFS File. 66
Message routing. 67
Using the Error Dump and Cleanup interface 68
Viewing the output data set 69
 Example of SMF Browser output 70

Chapter 6. Automation and recovery scenarios and guidelines 73

APPC automation and recovery scenarios 73
WLM automation and recovery scenarios 74
RACF automation and recovery scenarios 74
RRS automation and recovery scenarios 75
UNIX System Services automation and recovery scenarios 76
TCP/IP automation and recovery scenarios. 77
DB2 automation and recovery scenarios 77
CICS automation and recovery scenarios 78
IMS automation and recovery scenarios 79
LDAP automation and recovery scenarios 80
WebSphere Application Server for z/OS (Daemon) automation and recovery scenarios 81
Web server (servlet) automation and recovery scenarios 82

Chapter 7. Preparing for a call to IBM service. 83

Using the IPCS VERBEXIT subcommand to display diagnostic data 83
Setting trace controls for IBM service. 85
Setting dump controls for IBM service 88
Debugging Service details 89
 Startup 90
 JVM debug port 90
 JVM debug arguments. 90
 Debug class filters 90
 BSF debug port 90
 BSF logging level 90

Configuration problem settings	90
Configuration document validation	90
Cross validation	90
Configuration Problems	90
Scope	91
Message	91
Explanation	91
User action	91
Target Object	91
Severity	91
Local URI	91
Full URI	91
Validator classname	91

Chapter 8. Diagnosing and fixing problems: Resources for learning . . . 93

Chapter 9. Troubleshooting by task: What are you trying to do? 95

Troubleshooting installation problems	95
Error when installing WebSphere Application Server for z/OS	95
Installation completes but the administrative console does not start	95
Installation completes, but Sample applications do not work	96
Troubleshooting testing and first time run problems	97
A web resource does not display	97
Cannot access a data source	103
Cannot access an enterprise bean from a servlet, a JSP file, a stand-alone program, or another client	120
Cannot look up an object hosted by WebSphere Application Server from a servlet, JSP file, or other client	123
Access problems after enabling security	127
Errors after enabling security	130
Errors after configuring or enabling Secure Sockets Layer	134
Errors in messaging	137
Errors returned to a client sending a SOAP request	138
Errors connecting to WebSphere MQ and creating WebSphere MQ queue connection factory	139
Troubleshooting application runtime and management problems	139
Problems starting or using the wsadmin command	140
Web module or application server dies or hangs	142
Errors when trying to configure or enable security	143

Workload not getting distributed	144
Errors setting up multiserver environments	146
Cannot uninstall an application or remove a node or application server	149
Problems creating or using HTTP sessions	150
JSP source code shown by the Web server	152
Problems using tracing, logging or other troubleshooting features	153
Errors connecting to the administrative console from a Netscape browser	153
Resolving timeout conditions	156
Understanding how timers work	157
Guidelines for analyzing diagnostic data for timeout conditions	158
Identifying possible causes of and fixes for timeout conditions	160
Common timeout settings	162
Guidelines for altering timeout values	162
Debugging client exceptions	163
Debugging applications that hang	164
Debugging problems related to Java Message Service (JMS) support	166

Chapter 10. Troubleshooting by component: What is not working? . . . 169

Installation component troubleshooting tips	169
Administration and administrative console troubleshooting tips	169
Web Container troubleshooting tips	171
JDBC and data source troubleshooting tips	171
HTTP plug-in component troubleshooting tips	171
HTTP session manager troubleshooting tips	173
Naming services component troubleshooting tips	175
Messaging component troubleshooting tips	175
Universal Discovery, Description, and Integration, Web Service, and SOAP component troubleshooting tips	176
Enterprise bean and EJB container troubleshooting tips	176
Cannot restart the Deployment Manager monitoring policy	177
Security components troubleshooting tips	177
JSP processor troubleshooting tips	187
Object request broker component troubleshooting tips	188
Web services security troubleshooting tips	189

Notices 193

Trademarks and service marks 195

How to send your comments

Your feedback is important in helping to provide the most accurate and highest quality information.

- To send comments on articles in the WebSphere Application Server Information Center
 1. Display the article in your Web browser and scroll to the end of the article.
 2. Click on the **Feedback** link at the bottom of the article, and a separate window containing an e-mail form appears.
 3. Fill out the e-mail form as instructed, and click on **Submit feedback** .
- To send comments on PDF books, you can e-mail your comments to: **wasdoc@us.ibm.com** or fax them to 919-254-0206.

Be sure to include the document name and number, the WebSphere Application Server version you are using, and, if applicable, the specific page, table, or figure number on which you are commenting.

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.

Chapter 1. Troubleshooting or problem determination

As an activity, troubleshooting or problem determination encompasses a wide range of tasks that might need to be performed at any phase in product usage. In addition to built-in preventative measures, the product provides a variety of tools to make problem determination easier. . .

The troubleshooting section of the documentation helps you understand why your enterprise application, application server, or product installation is not working, and helps you resolve the problem.

There are several ways to find information for diagnosing and resolving problems:

- For tips on investigating common problems, organized by task, see Chapter 9, “Troubleshooting by task: What are you trying to do?,” on page 95.
- For help adding log and trace capability to your own application, see Chapter 3, “Adding logging and tracing to your application,” on page 7.

Chapter 2. Diagnosing and fixing problems

The purpose of this section is to aid you in understanding why your enterprise application, application server, or WebSphere Application Server is not working and to help you resolve the problem. 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.

1. For tips on investigating common problems organized according to tasks within WebSphere Application server, see Troubleshooting by task.
2. For tips on how to investigate common kinds of problems based on the component that is causing the problem, see Troubleshooting by component.
3. For help in using WebSphere Application Server utilities to help you diagnose the problem, see Working with diagnostic tools and controls.
4. For help in viewing diagnostic information like dumps, error logs and CTRACE information, see Viewing diagnostic information
5. The IBM Developer Kit and Runtime Environment, Java 2 Technology Edition, Version 1.4.1 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 <http://www.ibm.com/developerworks/java/jdk/diagnosis/>.
6. For current information available from IBM Support on known problems and their resolution, see the IBM Support page.
7. IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the IBM Support page.

Acquiring skills for problem determination

In a large-scale enterprise system such as the WebSphere Application Server for z/OS environment, diagnosis might require a variety of skills to progress from a symptom to fixing the underlying cause of that symptom. Because WebSphere Application Server for z/OS exploits many of the qualities and services that are unique to the z/OS operating system, diagnosing system-related problems might require skills in the following areas:

- Parallel sysplex
- TCP/IP
- Security Server (RACF) or the equivalent
- Database systems such as DB2 Universal Database for z/OS and OS/390
- UNIX Systems Services

You can find information for many of these topics in the publications available through the z/OS library Web site.

Similarly, diagnosing application-related problems might require a variety of skills because of the variety of application components that WebSphere Application Server for z/OS supports. Programmers who diagnose application problems in the WebSphere Application Server for z/OS environment need some familiarity with the following:

- The roles defined in the Sun Microsystems Java 2 Enterprise Edition Specification V1.3.

- The programming models and specifications for application components (Enterprise beans, Web applications, and client programs) supported in the J2EE 1.3 environment.
- The process of assembling, deploying, installing, and running server applications and clients in the WebSphere Application Server for z/OS environment.
- Various tools such as the WebSphere Application Server for z/OS error log, and the job logs for programs running on z/OS.

Working with diagnostic tools and controls

The following list summarizes the z/OS tools you will use to access and work with diagnostic information.

- **z/OS console**

The console displays configuration errors that cause the termination of the WebSphere Application Server for z/OS address spaces. Whatever goes to the console also goes to SYSLOG.

- **System log (SYSLOG)**

SYSLOG is the repository for all messages that have appeared on the operator console. It also contains warning and informational messages that might be helpful after a failure has occurred.

- **Job log**

The job log contains errors and warnings (non-termination) that are related to configuration. Anything that goes to the console and SYSLOG automatically goes to the job log.

- **System output (SYSOUT)**

SYSOUT is a batch log that usually contains diagnostic data from the Java Virtual Machine (JVM) that runs in the servant (region). Any messages written to stderr will end up in SYSOUT. In addition, SYSOUT might contain error messages that usually appear in the log stream, but were redirected to SYSOUT because the log stream was not available.

- **Error log**

The error log contains messages issued through JRas support, if any. In addition, the error log usually contains messages intended for IBM use only; these are messages that support actions, problems, or issues that are usually externalized through additional messages in other sources. When you work with IBM service, you might be asked to supply the error log so that service personnel can use these support messages to diagnose the problem.

Note: You must update the CFRM policy before using log streams that are CF-resident, such as the WebSphere error log and RRS logs. See Updating the CFRM policy for details.

- **SYSPRINT**

SYSPRINT contains component trace (CTRACE) output for clients, and for servants when WebSphere Application Server for z/OS is configured to use SYSPRINT rather than CTRACE buffers and data sets.

- **Component trace (CTRACE) data set**

CTRACE data sets contain diagnostic trace entries for various processes, depending on the trace options configured for WebSphere Application Server for z/OS.

- **Logrec**

When an error occurs, the system records information about the error in the logrec data set or the logrec log stream. The information provides you with a history of all hardware failures, selected software errors, and selected system conditions.

- **5.1+** **Transaction XA Partner Log**

This log is used for recovery of XA resources. When an application accesses XA resources, WebSphere Application Server stores information about the resource to enable XA transaction recovery. For instructions on how to use the Customization dialog to configure the Transaction XA Partner Log see Define Variables to configure base Application Server node - Definitions. For instructions on how to change the location of the Transaction XA Partner Log, see Transaction service settings.

To find additional information about these tools, and about the process of diagnosing problems on z/OS, use the z/OS Web Library to access the following books:

- *z/OS MVS Diagnosis: Procedures, GA22-7587* helps you diagnose problems in the MVS operating system, its subsystems, its components, and in applications running under the system.
- *z/OS MVS Diagnosis: Tools and Service Aids, GA22-7589* provides detailed information about tools and service aids that can help you diagnose problems. This book contains a guide on how to select the appropriate tool or service aid for your purposes, and also provides an overview of all the tools and service aids available.

Chapter 3. Adding logging and tracing to your application

Designers and developers of applications that run with or under WebSphere Application Server, such as servlets, JSP files, enterprise beans, client applications, and their supporting classes, may find it useful to use the same facility for generating messages that WebSphere Application Server itself uses, JRas.

This approach has advantages over simply adding `System.out.println()` statements to your code:

- Your messages appear in the WebSphere Application Server standard message format with additional data, such as a date and time stamp, added automatically.
- You can more easily correlate problems and events in your own application to problems and events associated with WebSphere Application Server components.
- You can take advantage of the WebSphere Application Server log file management features.
- You can view your messages with the WebSphere Application Server user-friendly Log Analyzer tool.

Unless you choose to extend the JRas framework, using the JRas API set is only slightly more complicated than writing basic `System.println()` statements.

Programming with the JRas framework

Use the JRas extensions to incorporate message logging and diagnostic trace into WebSphere Application Server applications. The JRas extensions are based on the stand-alone JRas logging toolkit.

1. Retrieve a reference to the JRas manager.
2. Retrieve message and trace loggers by using methods on the returned manager.
3. Call the appropriate methods on the returned message and trace loggers to create message and trace entries, as appropriate.

Understanding the JRas facility

Developing, deploying and maintaining applications are complex tasks. For example, when a running application encounters an unexpected condition it might not be able to complete a requested operation. In such a case you might want the application to inform the administrator that the operation has failed and give information as to why. This enables the administrator to take the proper corrective action. Application developers or maintainers might need to gather detailed information relating to the execution path of a running application in order to determine the root cause of a failure that is due to a code bug. The facilities that are used for these purposes are typically referred to as message logging and diagnostic trace.

Message logging (messages) and diagnostic trace (trace) are conceptually quite similar, but do have important differences. It is important for application developers to understand these differences in order to use these tools properly. To start with, the following operational definitions of messages and trace are provided.

Message

A message entry is an informational record intended to be viewed by end

users, systems administrators and support personnel. The text of the message must be clear, concise and interpretable by an end user. Messages are typically localized, meaning they are displayed in the national language of the end user. Although the destination and lifetime of messages might be configurable, some level of message logging is always enabled in normal system operation. Message logging must be used judiciously due to both performance considerations and the size of the message repository.

Trace A trace entry is an information record that is intended to be used by service engineers or developers. As such a trace record may be considerably more complex, verbose and detailed than a message entry. Localization support is typically not used for trace entries. Trace entries may be fairly inscrutable, understandable only by the appropriate developer or service personnel. It is assumed that trace entries are not written during normal runtime operation, but may be enabled as needed to gather diagnostic information.

WebSphere Application Server provides a message logging and diagnostic trace API that can be used by applications. This API is based on the stand-alone JRas logging toolkit which was developed by IBM. The stand-alone JRas logging toolkit is a collection of interfaces and classes that provide message logging and diagnostic trace primitives. These primitives are not tied to any particular product or platform. The stand-alone JRas logging toolkit provides a limited amount of support (typically referred to as systems management support), including log file configuration support based on property files.

As designed, the stand-alone JRas logging toolkit does not contain the support required for integration into the WebSphere Application Server runtime or for usage in a J2EE environment. To overcome these limitations, WebSphere Application Server provides a set of extension classes to address these shortcomings. This collection of extension classes is referred to as the JRas extensions. The JRas extensions do not modify the interfaces introduced by the stand-alone JRas logging toolkit, but simply provide the appropriate implementation classes. The conceptual structure introduced by the stand-alone JRas logging toolkit is described below. It is equally applicable to the JRas extensions.

JRas Concepts

The following is a basic overview of important concepts and constructs introduced by the stand-alone JRas logging toolkit. It is not meant to be an exhaustive overview of the capabilities of this logging toolkit, nor is it intended to be a detailed discussion of usage or programming paradigms. More detailed information, including code examples, is available in JRas extensions and its subtopics, including in the Javadoc for the various interfaces and classes that make up the logging toolkit.

Event Types

The stand-alone JRas logging toolkit defines a set of event types for messages and a set of event types for trace. Examples of message types include informational, warning and error. Examples of trace types include entry, exit and trace.

Event Classes

The stand-alone JRas logging toolkit defines both message and trace event classes.

Loggers

A logger is the primary object with which the user code interacts. Two

types of loggers are defined. These are message loggers and trace loggers. The set of methods on message loggers and trace loggers are different, since they provide different functionality. Message loggers create only message records and trace loggers create only trace records. Both types of loggers contain masks that indicates which categories of events the logger should process and which it should ignore. Although every JRas logger is defined to contain both a message and trace mask, the message logger only uses the message mask and the trace logger only uses the trace mask. For example, by setting a message logger's message mask to the appropriate state, it can be configured to process only Error messages and ignore Informational and Warning messages. Changing the state of a message logger's trace mask has no effect.

A logger contains one or more handlers to which it forwards events for further processing. When the user calls a method on the logger, the logger will compare the event type specified by the caller to its current mask value. If the specified type passes the mask check, the logger will create an event object to capture the information relating to the event that was passed to the logger method. This information may include information such as the names of the class and method which is logging the event, a message and parameters to log, among others. Once the logger has created the event object, it forwards the event to all handlers currently registered with the logger.

Methods that are used within the logging infrastructure itself should not make calls to the logger method. When an application uses an object that extends a thread class, implements the hashCode(), and makes a call to the logging infrastructure from that method, the result is a recursive loop.

Handlers

A handler provides an abstraction over an output device or event consumer. An example is a file handler, which knows how to write an event to a file. The handler also contains a mask that is used to further restrict the categories of events the handler will process. For example, a message logger may be configured to pass both warning and error events, but a handler attached to the message logger may be configured to only pass error events. Handlers also include formatters, which the handler invokes to format the data in the passed event before it is written to the output device.

Formatters

Handlers are configured with Formatters, which know how to format events of certain types. A handler may contain multiple formatters, each of which know how to format a specific class of event. The event object is passed to the appropriate formatter by the handler. The formatter returns formatted output to the handler, which then writes it to the output device.

JRas Extensions

JRas extensions

The stand-alone JRas logging toolkit defines interfaces and provides a variety of concrete classes that implement these interfaces. Since the stand-alone JRas logging toolkit was developed as a general purpose toolkit, the implementation classes do not contain the configuration interfaces and methods necessary for use in the WebSphere Application Server product. In addition, many of the implementation classes are not written appropriately for use in a J2EE environment. To overcome these shortcomings, WebSphere Application Server provides the appropriate

implementation classes that allow integration into the WebSphere Application Server environment. The collection of these implementation classes is referred to as the JRas extensions.

Usage Model

You can use the JRas extensions in three distinct operational modes:

Integrated

In this mode, message and trace records are written only to logs defined and maintained by the WebSphere Application Server runtime. This is the default mode of operation and is equivalent to the WebSphere Application Server 4.0 mode of operation.

stand-alone

In this mode, message and trace records are written solely to stand-alone logs defined and maintained by the user. You control which categories of events are written to which logs, and the format in which entries are written. You are responsible for configuration and maintenance of the logs. Message and trace entries are not written to WebSphere Application Server runtime logs.

Combined

In this mode message and trace records are written to both WebSphere Application Server runtime logs and to stand-alone logs that you must define, control, and maintain. You can use filtering controls to determine which categories of messages and trace are written to which logs.

The JRas extensions are specifically targeted to an integrated mode of operation. The integrated mode of operation can be appropriate for some usage scenarios, but there many scenarios are not adequately addressed by these extensions. Many usage scenarios require a stand-alone or combined mode of operation instead. A set of user extension points has been defined that allow the JRas extensions to be used in either a stand-alone or combined mode of operations.

JRas extension classes

WebSphere Application Server provides a base set of implementation classes that collectively are referred to as the JRas extensions. Many of these classes provide the appropriate implementations of loggers, handlers and formatters for use in a WebSphere Application Server environment. As previously noted, this collection of classes is targeted at an Integrated mode of operation. If you choose to use the JRas extensions in either stand-alone or combined mode, you can reuse the logger and manager class provided by the extensions, but you must provide your own implementations of handlers and formatters.

WebSphere Application Server Message and Trace loggers

The message and trace loggers provided by the stand-alone JRas logging toolkit cannot be directly used in the WebSphere Application Server environment. The JRas extensions provide the appropriate logger implementation classes. Instances of these message and trace logger classes are obtained directly and exclusively from the WebSphere Application Server Manager class, described below. You cannot directly instantiate message and trace loggers. Obtaining loggers in any manner other than directly from the Manager is not allowed. Doing so is a direct violation of the programming model.

The message and trace loggers instances obtained from the WebSphere Application Server Manager class are subclasses of the `RASMessageLogger()` and

RASTraceLogger() classes provided by the stand-alone JRes logging toolkit. The RASMessageLogger() and RASTraceLogger() classes define the set of methods that are directly available. Public methods introduced by the JRes extensions logger subclasses cannot be called directly by user code. Doing so is a violation of the programming model.

Loggers are named objects and are identified by name. When the Manager class is called to obtain a logger, the caller is required to specify a name for the logger. The Manager class maintains a name-to-logger instance mapping. Only one instance of a named logger will ever be created within the lifetime of a process. The first call to the Manager with a particular name will result in the logger being created and configured by the Manager. The Manager will cache a reference to the instance, then return it to the caller. Subsequent calls to the Manager that specify the same name will result in a reference to the cached logger being returned. Separate namespaces are maintained for message and trace loggers. This means a single name can be used to obtain both a message logger and a trace logger from the Manager, without ambiguity, and without causing a namespace collision.

In general, loggers have no predefined granularity or scope. A single logger could be used to instrument an entire application. Or users may determine that having a logger per class is more desirable. Or the appropriate granularity may lie somewhere in between. Partitioning an application into logging domains is rightfully determined by the application writer.

The WebSphere Application Server logger classes obtained from the Manager are thread-safe. Although the loggers provided as part of the stand-alone JRes logging toolkit implement the serializable interface, in fact loggers are not serializable. Loggers are stateful objects, tied to a Java virtual machine instance and are not serializable. Attempting to serialize a logger is a violation of the programming model.

Please note that there is no provision for allowing users to provide their own logger subclasses for use in a WebSphere Application Server environment.

WebSphere Application Server handlers

WebSphere Application Server provides the appropriate handler class that is used to write message and trace events to the WebSphere Application Server run-time logs. You cannot configure the WebSphere Application Server handler to write to any other destination. The creation of a WebSphere Application Server handler is a restricted operation and not available to user code. Every logger obtained from the Manager comes preconfigured with an instance of this handler already installed. You can remove the WebSphere Application Server handler from a logger when you want to run in stand-alone mode. Once you have removed it, you cannot add the WebSphere Application Server handler again to the logger from which it was removed (or any other logger). Also, you cannot directly call any method on the WebSphere Application Server handler. Attempting to create an instance of the WebSphere Application Server handler, to call methods on the WebSphere Application Server handler or to add a WebSphere Application Server handler to a logger by user code is a violation of the programming model.

WebSphere Application Server formatters

The WebSphere Application Server handler comes preconfigured with the appropriate formatter for data that is written to WebSphere Application Server logs. The creation of a WebSphere Application Server formatter is a restricted

operation and not available to user code. No mechanism exists that allows the user to obtain a reference to a formatter installed in a WebSphere Application Server handler, or to change the formatter a WebSphere Application Server handler is configured to use.

WebSphere Application Server manager

WebSphere Application Server provides a Manager class located in the `com.ibm.websphere.ras` package. All message and trace loggers must be obtained from this Manager. A reference to the Manager is obtained by calling the static `Manager.getManager()` method. Message loggers are obtained by calling the `createRASMessageLogger()` method on the Manager. Trace loggers are obtained by calling the `createRASTraceLogger()` method on the Manager class.

The manager also supports a *group* abstraction that is useful when dealing with trace loggers. The group abstraction allows multiple, unrelated trace loggers to be registered as part of a named entity called a group. WebSphere Application Server provides the appropriate systems management facilities to manipulate the trace setting of a group, similar to the way the trace settings of an individual trace logger.

For example, suppose component A consist of 10 classes. Suppose each class is configured to use a separate trace logger. Suppose all 10 trace loggers in the component are registered as members of the same group (for example `Component_A_Group`). You can then turn on trace for a single class. Or you can turn on trace for all 10 classes in a single operation using the group name if you want a component trace. Group names are maintained within the namespace for trace loggers.

Extending the JRas framework

Since the JRas extensions classes do not provide the flexibility and behavior required for many scenarios, a variety of extension points have been defined. You are allowed to write your own implementation classes to obtain the required behavior.

In general, the JRas extensions require you to call the Manager class to obtain a message logger or trace logger. No provision is made to allow you to provide your own message or trace logger subclasses. In general, user-provided extensions cannot be used to affect the integrated mode of operation. The behavior of the integrated mode of operation is solely determined by the WebSphere Application Server run-time and the JRas extensions classes.

Handlers

The stand-alone JRas logging toolkit defines the `RASHandler` interface. All handlers must implement this interface. You can write your own handler classes that implement the `RASHandler` interface. You should directly create instances of user-defined handlers and add them to the loggers obtained from the Manager.

The stand-alone JRas logging toolkit provides several handler implementation classes. These handler classes are inappropriate for usage in the J2EE environment. You cannot directly use or subclass any of the Handler classes provided by the stand-alone JRas logging toolkit. Doing so is a violation of the programming model.

Formatters

The stand-alone JRas logging toolkit defines the `RASIFormatter` interface. All formatters must implement this interface. You can write your own formatter classes that implement the `RASIFormatter` interface. You can only add these classes to a user-defined handler. WebSphere Application Server handlers cannot be configured to use user-defined formatters. Instead, directly create instances of your formatters and add them to the your handlers appropriately.

As with handlers, the stand-alone JRas logging toolkit provides several formatter implementation classes. Direct usage of these formatter classes is not supported.

Message event types

The stand-alone JRas toolkit defines message event types in the `RASIMessageEvent` interface. In addition, the WebSphere Application Server reserves a range of message event types for future use. The `RASIMessageEvent` interface defines three types, with values of `0x01`, `0x02`, and `0x04`. The values `0x08` through `0x8000` are reserved for future use. You can provide your own message event types by extending this interface appropriately. User-defined message types must have a value of `0x1000` or greater.

Message loggers retrieved from the Manager have their message masks set to *pass* or process all message event types defined in the `RASIMessageEvent` interface. In order to process user-defined message types, you must manually set the message logger mask to the appropriate state by user code after the message logger has been obtained from the Manager. WebSphere Application Server does not provide any built-in systems management support for managing any message types.

Message event objects

The stand-alone JRas toolkit provides a `RASMessageEvent` implementation class. When a message logging method is called on the message logger, and the message type is currently enabled, the logger creates and distributes an event of this class to all handlers currently registered with that logger.

You can provide your own message event classes, but they must implement the `RASIEvent` interface. You must directly create instances of such user-defined message event classes. Once it is created, pass your message event to the message logger by calling the message logger's `fireRASEvent()` method directly. WebSphere Application Server message loggers cannot directly create instances of user-defined types in response to calling a logging method (`msg()`, `message()`...) on the logger. In addition, instances of user-defined message types are never processed by the WebSphere Application Server handler. You cannot create instances of the `RASMessageEvent` class directly.

Trace event types

The stand-alone JRas toolkit defines trace event types in the `RASITraceEvent` interface. You can provide your own trace event types by extending this interface appropriately. In such a case you must ensure that the values for the user-defined trace event types do not collide with the values of the types defined in the `RASITraceEvent` interface.

Trace loggers retrieved from the Manager typically have their trace masks set to reject all types. A different starting state can be specified by using WebSphere

Application Server systems management facilities. In addition, the state of the trace mask for a logger can be changed at run-time using WebSphere Application Server systems management facilities.

In order to process user-defined trace types, the trace logger mask must be manually set to the appropriate state by user code. WebSphere Application Server systems management facilities cannot be used to manage user-defined trace types, either at start time or run-time.

Trace event objects

The stand-alone JRas toolkit provides a `RASTraceEvent` implementation class. When a trace logging method is called on the WebSphere Application Server trace logger and the type is currently enabled, the logger creates and distributes an event of this class to all handlers currently registered with that logger.

You can provide your own trace event classes. Such trace event classes must implement the `RASIEvent` interface. You must create instances of such user-defined event classes directly. Once it is created, pass the trace event to the trace logger by calling the trace logger's `fireRASEvent()` method directly. WebSphere Application Server trace loggers cannot directly create instances of user-defined types in response to calling a trace method (`entry()`, `exit()`, `trace()`) on the trace logger. In addition, instances of user-defined trace types are never processed by the WebSphere Application Server handler. You cannot create instances of the `RASTraceEvent` class directly.

User defined types, user defined events and WebSphere Application Server

By definition, the WebSphere Application Server handler will process user-defined message or trace types, or user-defined message or trace event classes. Message and trace entries of either a user-defined type or user-defined event class cannot be written to the WebSphere Application Server run-time logs.

Writing User Extensions: General Considerations

You can configure the WebSphere Application Server to use Java 2 security to restrict access to protected resources such as the file system and sockets. Since user written extensions typically access such protected resources, user written extensions must contain the appropriate security checking calls, using `AccessController.doPrivileged()` calls. In addition, the user written extensions must contain the appropriate policy file. In general, it is recommended that you locate user written extensions in a separate package. It is your responsibility to restrict access to the user written extensions appropriately.

Writing a handler

User written handlers must implement the `RASIHandler` interface. The `RASIHandler` interface extends the `RASIMaskChangeGenerator` interface, which extends the `RASIObjecT` interface. A short discussion of the methods introduced by each of these interfaces follows, along with implementation pointers. For more in depth information on any of the particular interfaces or methods, see the corresponding product Javadoc.

RASIObjecT interface

The RASIObjecT interface is the base interface for stand-alone JRas logging toolkit classes that are stateful or configurable, such as loggers, handlers and formatters.

- The stand-alone JRas logging toolkit supports rudimentary properties-file based configuration. To implement this configuration support, the configuration state is stored as a set of key-value pairs in a properties file. The methods *public Hashtable getConfig()* and *public void setConfig(Hashtable ht)* are used to get and set the configuration state. The JRas extensions do not support properties based configuration and it is recommended that these methods be implemented as no-operations. You can implement your own properties based configuration using these methods.
- Loggers, handlers and formatters can be named objects. For example, the JRas extensions require the user to provide a name for the loggers that are retrieved from the manager. You can name your handlers. The methods *public String getName()* and *public void setName(String name)* are provided to get or set the name field. The JRas extensions currently do not call these methods on user handlers. You can implement these methods as you want, including as no operations.
- Loggers, handlers and formatters can also contain a description field. The methods *public String getDescription()* and *public void setDescription(String desc)* can be used to get or set the description field. The JRas extensions currently do not use the description field. You can implement these methods as you want, including as no operations.
- The method *public String getGroup()* is provided for usage by the RASManager. Since the JRas extensions provide their own Manager class, this method is never called. It is recommended you implement this as a no-operation.

RASIMaskChangeGenerator interface

The RASIMaskChangeGenerator interface is the interface that defines the implementation methods for filtering of events based on a mask state. This means that it is currently implemented by both loggers and handlers. By definition, an object that implements this interface contains both a message mask and a trace mask, although both need not be used. For example, message loggers contain a trace mask, but the trace mask is never used since the message logger never generates trace events. Handlers however can actively use both mask values. For example a single handler could handle both message and trace events.

- The methods *public long getMessageMask()* and *public void setMessageMask(long mask)* are used to get or set the value of the message mask. The methods *public long getTraceMask()* and *public void setTraceMask(long mask)* are used to get or set the value of the trace mask.

In addition, this interface introduces the concept of *calling back* to interested parties when a mask changes state. The callback object must implement the RASIMaskChangeListener interface.

- The methods *public void addMaskChangeListener(RASIMaskChangeListener listener)* and *public void removeMaskChangeListener(RASIMaskChangeListener listener)* are used to add or remove listeners to the handler. The method *public Enumeration getMaskChangeListeners()* returns an Enumeration over the list of currently registered listeners. The method *public void fireMaskChangedEvent(RASMaskChangeEvent mc)* is used to call back all the registered listeners to inform them of a mask change event.

For efficiency reasons, the JRas extensions message and trace loggers implement the RASIMaskChangeListener interface. The logger implementations maintain a

"composite mask" in addition to the logger's own mask. The logger's composite mask is formed by logically *or'ing* the appropriate masks of all handlers that are registered to that logger, then *and'ing* the result with the logger's own mask. For example, the message logger's composite mask is formed by *or'ing* the message masks of all handlers registered with that logger, then *and'ing* the result with the logger's own message mask.

This means that all handlers are required to properly implement these methods. In addition, when a user handler is instantiated, the logger it is to be added to should be registered with the handler using the *addMaskChangeListener()* method. When either the message mask or trace mask of the handler is changed, the logger must be called back to inform it of the mask change. This allows the logger to dynamically maintain the composite mask.

The *RASMaskChangedEvent* class is defined by the stand-alone JRas logging toolkit. Direct usage of that class by user code is allowed in this context.

In addition the *RASIMaskChangeGenerator* introduces the concept of caching the names of all message and trace event classes that the implementing object will process. The intent of these methods is to allow a management program such as a GUI to retrieve the list of names, introspect the classes to determine the event types that they might possibly process and display the results. The JRas extensions do not ever call these methods, so they can be implemented as no operations, if desired.

- The methods *public void addMessageEventClass(String name)* and *public void removeMessageEventClass(String name)* can be called to add or remove a message event class name from the list. The method *public Enumeration getMessageEventClasses()* will return an enumeration over the list of message event class names. Similarly, the *public void addTraceEventClass(String name)* and *public void removeTraceEventClass(String name)* can be called to add or remove a trace event class name from the list. The method *public Enumeration getTraceEventClasses()* will return an enumeration over the list of trace event class names.

RASHandler interface

The *RASHandler* interface introduces the methods that are specific to the behavior of a handler.

The *RASHandler* interface as provided by the stand-alone JRas logging toolkit supports handlers that run in either a synchronous or asynchronous mode. In asynchronous mode, events are typically queued by the calling thread and then written by a worker thread. Since spawning of threads is not allowed in the WebSphere Application Server environment, it is expected that handlers will not queue or batch events, although this is not expressly prohibited.

- The methods *public int getMaximumQueueSize()* and *public void setMaximumQueueSize(int size)* throw *IllegalStateException* are provided to manage the maximum queue size. The method *public int getQueueSize()* is provided to query the actual queue size.
- The methods *public int getRetryInterval()* and *public void setRetryInterval(int interval)* support the notion of error retry, which again implies some type of queueing.
- The methods *public void addFormatter(RASFormatter formatter)*, *public void removeFormatter(RASFormatter formatter)* and *public Enumeration*

getFormatters() are provided to manage the list of formatters that the handler can be configured with. Different formatters can be provided for different event classes, if appropriate.

- The methods *public void openDevice()*, *public void closeDevice()* and *public void stop()* are provided to manage the underlying device that the handler abstracts.
- The methods *public void logEvent(RASIEvent event)* and *public void writeEvent(RASIEvent event)* are provided to actually pass events to the handler for processing.

Writing a formatter

User written formatters must implement the RASIFormatter interface. The RASIFormatter interface extends the RASIObjcet interface. The implementation of the RASIObjcet interface is the same for both handlers and formatters. A short discussion of the methods introduced by the RASIFormatter interface follows. For more in depth information on the methods introduced by this interface, see the corresponding product javadoc.

RASIFormatter interface

- The methods *public void setDefault(boolean flag)* and *public boolean isDefault()* are used by the concrete RASHandler classes provided by the stand-alone JRas logging toolkit to determine if a particular formatter is the default formatter. Since these RASHandler classes must never be used in a WebSphere Application Server environment, the semantic significance of these methods can be determined by the user.
- The methods *public void addEventClass(String name)*, *public void removeEventClass(String name)* and *public Enumeration getEventClasses()* are provided to determine which event classes a formatter can be used to format. You can provide the appropriate implementations as you see fit.
- The method *public String format(RASIEvent event)* is called by handler objects and returns a formatted String representation of the event.

Programming model summary

The programming model described in this section builds upon and summarizes some of the concepts already introduced. This section also formalizes usage requirements and restrictions. Use of the WebSphere Application Server JRas extensions in a manner that does not conform to the following programming guidelines is prohibited.

As described previously, you can use the WebSphere Application Server JRas extensions in three distinct operational modes. The programming models concepts and restrictions apply equally across all modes of operation.

- You must not use implementation classes provided by the stand-alone JRas logging toolkit directly, unless specifically noted otherwise. Direct usage of those classes is not supported. IBM Support will provide no diagnostic aid or bug fixes relating to direct usage of classes provided by the stand-alone JRas logging toolkit.
- You must obtain message and trace loggers directly from the Manager class. You cannot directly instantiate loggers.
- There is no provision that allows you to replace the WebSphere Application Server message and trace logger classes.
- You must guarantee that the logger names passed to the Manager are unique, and follow the naming constraints documented below. Once a logger is obtained from the Manager, you must not attempt to change the name of the logger by calling the setName() method.

- Named loggers can be used more than once. For any given name, the first call to the Manager results in the Manager creating a logger that is associated with that name. Subsequent calls to the Manager that specify the same name result in a reference to the existing logger being returned.
- The Manager maintains a hierarchical namespace for loggers. It is recommended but not required that a dot-separated, fully qualified class name be used to identify any given logger. Other than dots or periods, logger names cannot contain any punctuation characters, such as asterisk (*), comma(.), equals sign(=), colon(:), or quotes.
- Group names must comply with the same naming restrictions as logger names.
- The loggers returned from the Manager are subclasses of the `RASMessageLogger` and `RASTraceLogger` provided by the stand-alone JRas logging toolkit. You are allowed to call any public method defined by the `RASMessageLogger` and `RASTraceLogger` classes. You are not allowed to call any public method introduced by the provided subclasses.
- If you want to operate in either stand-alone or combined mode, you must provide your own Handler and Formatter subclasses. You are not allowed to use the Handler and Formatter classes provided by the stand-alone JRas logging toolkit. User written Handlers and Formatters must conform to the documented guidelines.
- Loggers obtained from the Manager come with a WebSphere Application Server handler installed. This handler will write message and trace records to logs defined by the WebSphere Application Server runtime. Manage these logs using the provided systems management interfaces.
- You can programmatically add and remove user-defined Handlers from a logger at any time. Multiple additions and removals of user defined handlers are allowed. You are responsible for creating an instance of the handler to add, configuring the handler by setting the handler's mask value and formatter appropriately, then adding the handler to the logger using the `addHandler()` method. You are responsible for programmatically updating the masks of user-defined handlers as appropriate.
- You may get a reference to the handler installed within a logger by calling the `getHandlers()` method on the logger and processing the results. You must not call any methods on the handler obtained in this fashion. You are allowed to remove the WebSphere Application Server handler from the logger by calling the logger's `removeHandler()` method, passing in the reference to the WebSphere Application Server handler. Once removed, the WebSphere Application Server handler cannot be re-added to the logger.
- You are allowed to define your own message type. The behavior of user-defined message types and restrictions on their definitions is discussed in Extending the JRas framework.
- You are allowed to define your own message event classes. The usage of user-defined message event classes is discussed in Extending the JRas framework.
- You are allowed to define your own trace types. The behavior of user-defined trace types and restrictions on your definitions is discussed in Extending the JRas framework.
- You are allowed to define your own trace event classes. The usage of user-defined trace event classes is discussed in Extending the JRas framework.
- You must programmatically maintain the bits in the message and trace logger masks that correspond to any user-defined types. If WebSphere Application Server facilities are being used to manage the predefined types, these updates must not modify the state of any of the bits corresponding to those types. If you are assuming ownership responsibility for the predefined types then you can change all bits of the masks.

JRas Messages and Trace event types

This section describes JRas message and trace event types.

Event types

The base message and trace event types defined by the stand-alone JRas logging toolkit are not the same as the "native" types recognized by the WebSphere Application Server run-time. Instead the basic JRas types are mapped onto the native types. This mapping may vary by platform or edition. The mapping is discussed below.

Platform Message Event Types

The message event types that are recognized and processed by the WebSphere Application Server runtime are defined in the `RASIMessageEvent` interface provided by the stand-alone JRas logging toolkit. These message types are mapped onto the native message types as follows.

WebSphere Application Server native type	JRas RASIMessageEvent type
Audit	TYPE_INFO, TYPE_INFORMATION
Warning	TYPE_WARN, TYPE_WARNING
Error	TYPE_ERR, TYPE_ERROR

Platform Trace Event Types

The trace event types recognized and processed by the WebSphere Application Server runtime are defined in the `RASITraceEvent` interface provided by the stand-alone JRas logging toolkit. The `RASITraceEvent` interface provides a rich and overly complex set of types. This interface defines both a simple set of levels, as well as a set of enumerated types.

- For a user who prefers a simple set of levels, `RASITraceEvent` provides `TYPE_LEVEL1`, `TYPE_LEVEL2`, and `TYPE_LEVEL3`. The implementations provide support for this set of levels. The levels are hierarchical (that is, enabling level 2 will also enable level 1, enabling level 3 also enables levels 1 and 2).
- For users who prefer a more complex set of values that can be *OR'd* together, `RASITraceEvent` provides `TYPE_API`, `TYPE_CALLBACK`, `TYPE_ENTRY_EXIT`, `TYPE_ERROR_EXC`, `TYPE_MISC_DATA`, `TYPE_OBJ_CREATE`, `TYPE_OBJ_DELETE`, `TYPE_PRIVATE`, `TYPE_PUBLIC`, `TYPE_STATIC`, and `TYPE_SVC`.

The trace event types are mapped onto the native trace types as follows:

Mapping WebSphere Application Server trace types to JRas `RASITraceEvent` "Level" types.

WebSphere Application Server native type	JRas RASITraceEvent level type
Event	TYPE_LEVEL1
EntryExit	TYPE_LEVEL2
Debug	TYPE_LEVEL3

Mapping WebSphere Application Server trace types to JRas `RASITraceEvent` enumerated types.

WebSphere Application Server native type	JRas RASITraceEvent enumerated types
Event	TYPE_ERROR_EXC, TYPE_SVC, TYPE_OBJ_CREATE, TYPE_OBJ_DELETE
EntryExit	TYPE_ENTRY_EXIT, TYPE_API, TYPE_CALLBACK, TYPE_PRIVATE, TYPE_PUBLIC, TYPE_STATIC
Debug	TYPE_MISC_DATA

For simplicity, it is recommended that one or the other of the tracing type methodologies is used consistently throughout the application. For users who decide to use the non-level types, it is further recommended that you choose one type from each category and use those consistently throughout the application to avoid confusion.

Message and Trace parameters

The various message logging and trace method signatures accept parameter types of Object, Object[] and Throwable. WebSphere Application Server will process and format the various parameter types as follows.

Primitives

Primitives, such as int and long are not recognized as subclasses of Object and cannot be directly passed to one of these methods. A primitive value must be transformed to a proper Object type (Integer, Long) before being passed as a parameter.

Object toString() is called on the object and the resulting String is displayed. The toString() method should be implemented appropriately for any object passed to a message logging or trace method. It is the responsibility of the caller to guarantee that the toString() method does not display confidential data such as passwords in clear text, and does not cause infinite recursion.

Object[]

The Object[] is provided for the case when more than one parameter is passed to a message logging or trace method. toString() is called on each Object in the array. Nested arrays are not handled. (i.e. none of the elements in the Object array should be an array).

Throwable

The stack trace of the Throwable is retrieved and displayed.

Array of Primitives

An array of primitive (e.g. byte[], int[] is recognized as an Object, but is treated somewhat as a second cousin of Object by Java code. In general, arrays of primitives should be avoided, if possible. If arrays of primitives are passed, the results are indeterminate and may change depending on the type of array passed, the API used to pass the array and the release of the product. For consistent results, user code should preprocess and format the primitive array into some type of String form before passing it to the method. If such preprocessing is not performed, the following may result.

- [B@924586a0b - This is deciphered as "a byte array at location X". This is typically returned when an array is passed as a member of an Object[]. It is the result of calling toString() on the byte[].
- Illegal trace argument : array of long. This is typically returned when an array of primitives is passed to a method taking an Object.
- 01040703... : the hex representation of an array of bytes. Typically this may be seen when a byte array is passed to a method taking a single Object. This behavior is subject to change and should not be relied on.

- "1" "2" ... : The String representation of the members of an int[] formed by converting each element to an Integer and calling toString on the Integers. This behavior is subject to change and should not be relied on.
- [Ljava.lang.Object;@9136fa0b : An array of objects. Typically this is seen when an array containing nested arrays is passed.

Controlling message logging

Writing a message to a WebSphere Application Server log requires that the message type passes three levels of filtering or screening.

1. The message event type must be one of the message event types defined in the RASIMessageEvent interface.
2. Logging of that message event type must be enabled by the state of the message logger's mask.
3. The message event type must pass any filtering criteria established by the WebSphere Application Server run-time itself.

When a WebSphere Application Server logger is obtained from the Manager, the initial setting of the mask is to forward all native message event types to the WebSphere Application Server handler. It is possible to control what messages get logged by programmatically setting the state of the message logger's mask.

Some editions of the product allow the user to specify a message filter level for a server process. When such a filter level is set, only messages at the specified severity levels are written to WebSphere Application Server logs. This means that messages types that pass the message logger's mask check may be filtered out by the WebSphere Application Server itself.

Controlling Tracing

Each edition of the product provides a mechanism for enabling or disabling trace. The various editions may support static trace enablement (trace settings are specified before the server is started), dynamic trace enablement (trace settings for a running server process can be dynamically modified) or both.

Writing a trace record to a WebSphere Application Server requires that the trace type passes three levels of filtering or screening.

1. The trace event type must be one of the trace event types defined in the RASITraceEvent interface.
2. Logging of that trace event type must be enabled by the state of the trace logger's mask.
3. The trace event type must pass any filtering criteria established by the WebSphere Application Server run-time itself.

When a logger is obtained from the Manager, the initial setting of the mask is to suppress all trace types. The exception to this rule is the case where the WebSphere Application Server run-time supports static trace enablement and a non-default startup trace state for that trace logger has been specified. Unlike message loggers, the WebSphere Application Server may dynamically modify the state of a trace loggers trace mask. WebSphere Application Server will only modify the portion of the trace logger's mask corresponding to the values defined in the RASITraceEvent interface. WebSphere Application Server will not modify undefined bits of the mask that may be in use for user defined types.

When the dynamic trace enablement feature available on some platforms is used, the trace state change is reflected both in the Application Server run-time and the

trace loggers trace mask. If user code programmatically changes the bits in the trace mask corresponding to the values defined by in the RASITraceEvent interface, the trace logger's mask state and the run-time state will become unsynchronized and unexpected results will occur. Therefore, programmatically changing the bits of the mask corresponding to the values defined in the RASITraceEvent interface is not allowed.

Instrumenting an application with JRas extensions

To instrument an application using the WebSphere Application Server JRas extensions, perform the following steps:

1. Determine the mode the extensions will be used in: integrated, stand-alone or combined.
2. If the extensions will be used in either stand-alone or combined mode, create the necessary handler and formatter classes.
3. If localized messages will be used by the application, create a resource bundle as described in Creating JRas resource bundles and message files.
4. In the application code, get a reference to the Manager class and create the manager and logger instances as described in Creating JRas manager and logger instances.
5. Insert the appropriate message and trace logging statements in the application as described in Creating JRas manager and logger instances.

Creating JRas resource bundles and message files

The WebSphere Application Server message logger provides the `message()` and `msg()` methods to allow the user to log localized messages. In addition, it provides the `textMessage()` method for logging of messages that are not localized. Applications can use either or both, as appropriate.

The mechanism for providing localized messages is the Resource Bundle support provided by the IBM Developer Java Technology Edition. If you are not familiar with resource bundles as implemented by the Developer's Kit, you can get more information from various texts, or by reading the Javadoc for the `java.util.ResourceBundle`, `java.util.ListResourceBundle` and `java.util.PropertyResourceBundle` classes, as well as the `java.text.MessageFormat` class.

The `PropertyResourceBundle` is the preferred mechanism to use. In addition, note that the JRas extensions do not support the extended formatting options such as `{1, date}` or `{0,number, integer}` that are provided by the `MessageFormat` class.

You can forward messages that are written to the internal WebSphere Application Server logs to other processes for display. For example, messages displayed on the administrator console, which can be running in a different location than the server process, can be localized using the *late binding* process. Late binding means that WebSphere Application Server does not localize messages when they are logged, but defers localization to the process that displays the message.

To properly localize the message, the displaying process must have access to the resource bundle where the message text is stored. This means that you must package the resource bundle separately from the application, and install it in a location where the viewing process can access it. If you do not want to take these steps, you can use the early binding technique to localize messages as they are logged.

The two techniques are described as follows:

Early binding

The application must localize the message before logging it. The application looks up the localized text in the resource bundle and formats the message. When formatting is complete, the application logs the message using the `textMessage()` method. Use this technique to package the application's resource bundles with the application.

Late binding

The application can choose to have the WebSphere Application Server runtime localize the message in the process where it is displayed. Using this technique, the resource bundles are packaged in a stand-alone `.jar` file, separately from the application. You must then install the resource bundle `.jar` file on every machine in the installation from which an administrator's console or log viewing program might be run. You must install the `.jar` file in a directory that is part of the extensions classpath. In addition, if you forward logs to IBM service, you must also forward the `.jar` file containing the resource bundles.

To create a resource bundle, perform the following steps.

1. Create a text properties file that lists message keys and the corresponding messages. The properties file must have the following characteristics:
 - Each property in the file is terminated with a line-termination character.
 - If a line contains only white space, or if the first non-white space character of the line is the pound sign symbol (`#`) or exclamation mark (`!`), the line is ignored. The `#` and `!` characters can therefore be used to put comments into the file.
 - Each line in the file, unless it is a comment or consists only of white space, denotes a single property. A backslash (`\`) is treated as the line-continuation character.
 - The syntax for a property file consists of a key, a separator, and an element. Valid separators include the equal sign (`=`), colon (`:`), and white space ().
 - The key consists of all characters on the line from the first non-white space character to the first separator. Separator characters can be included in the key by escaping them with a backslash (`\`), but doing this is not recommended, because escaping characters is error prone and confusing. It is instead recommended that you use a valid separator character that does not appear in any keys in the properties file.
 - White space after the key and separator is ignored until the first non-white space character is encountered. All characters remaining before the line-termination character define the element.

See the Java documentation for the `java.util.Properties` class for a full description of the syntax and construction of properties files.

2. The file can then be translated into localized versions of the file with language-specific file names (for example, a file named `DefaultMessages.properties` can be translated into `DefaultMessages_de.properties` for German and `DefaultMessages_ja.properties` for Japanese).
3. When the translated resource bundles are available, write them to a system-managed persistent storage medium. Resource bundles are then used to convert the messages into the requested national language and locale.
4. When a message logger is obtained from the Jiras manager, it can be configured to use a particular resource bundle. Messages logged via the `message()` API will use this resource bundle when message localization is performed. At run time, the user's locale setting is used to determine the properties file from

which to extract the message specified by a message key, thus ensuring that the message is delivered in the correct language.

5. If the message loggers `msg()` method is called, a resource bundle name must be explicitly provided.

The application locates the resource bundle based on the file's location relative to any directory in the classpath. For instance, if the property resource bundle named `DefaultMessages.properties` is located in the `baseDir/subDir1/subDir2/resources` directory and `baseDir` is in the class path, the name `subDir1.subDir2.resources.DefaultMessage` is passed to the message logger to identify the resource bundle.

Developing JRas resource bundles: Resource bundle sample

You can create resource bundles in several ways. The best and easiest way is to create a properties file that supports a `PropertiesResourceBundle`. This sample shows how to create such a properties file.

For this sample, four localizable messages are provided. The properties file is created and the key-value pairs inserted into it. All the normal properties files conventions and rules apply to this file. In addition, the creator must be aware of other restrictions imposed on the values by the Java `MessageFormat` class. For example, apostrophes must be "escaped" or they will cause a problem. Also avoid use of non-portable characters. WebSphere Application Server does not support usage of extended formatting conventions that the `MessageFormat` class supports, such as `{1, date}` or `{0,number, integer}`.

Assume that the base directory for the application that uses this resource bundle is "baseDir" and that this directory will be in the classpath. Assume that the properties file is stored in a subdirectory of `baseDir` that is not in the classpath (e.g. `baseDir/subDir1/subDir2/resources`). In order to allow the messages file to be resolved, the name `subDir1.subDir2.resources.DefaultMessage` is used to identify the `PropertyResourceBundle` and is passed to the message logger.

For this sample, the properties file is named `DefaultMessages.properties`.

```
# Contents of DefaultMessages.properties file
MSG_KEY_00=A message with no substitution parameters.
MSG_KEY_01=A message with one substitution parameter: parm1={0}
MSG_KEY_02=A message with two substitution parameters: parm1={0}, parm2 = {1}
MSG_KEY_03=A message with three parameter: parm1={0}, parm2 = {1}, parm3={2}
```

Once the file `DefaultMessages.properties` is created, the file can be sent to a translation center where the localized versions will be generated.

Creating JRas manager and logger instances

You can use the JRas extensions in integrated, stand-alone, or combined mode. Configuration of the application will vary depending on the mode of operation, but usage of the loggers to log message or trace entries is identical in all modes of operation.

Integrated mode is the default mode of operation. In this mode, message and trace events are sent to the WebSphere Application Server logs. See [Setting up for integrated JRas operation](#) for information on configuring for this mode of operation.

In the combined mode, message and trace events are logged to both WebSphere Application Server and user-defined logs. See *Setting up for combined JRas operation* for more information on configuring for this mode of operation.

In the stand-alone mode, message and trace events are logged only to user-defined logs. See *Setting up for stand-alone JRas operation* for more information on configuring for this mode of operation.

Using the message and trace loggers

Regardless of the mode of operation, the use of message and trace loggers is the same. See *Creating JRas resource bundles and message files* for more information on using message and trace loggers.

Using a message logger

The message logger is configured to use the `DefaultMessages` resource bundle. Message keys must be passed to the message loggers if the loggers are using the `message()` API.

```
msgLogger.message(RASIMessageEvent.TYPE_WARNING, this,
    methodName, "MSG_KEY_00");
... msgLogger.message(RASIMessageEvent.TYPE_WARN, this,
    methodName, "MSG_KEY_01", "some string");
```

If message loggers use the `msg()` API, you can specify a new resource bundle name.

```
msgLogger.msg(RASIMessageEvent.TYPE_ERR, this, methodName,
    "ALT_MSG_KEY_00", "alternateMessageFile");
```

You can also log a text message. If you are using the `textMessage` API, no message formatting is done.

```
msgLogger.textMessage(RASIMessageEvent.TYPE_INFO, this, methodName, "String and Integer",
    "A String", new Integer(5));
```

Using a trace logger

Since trace is normally disabled, trace methods should be guarded for performance reasons.

```
private void methodX(int x, String y, Foo z)
{
    // trace an entry point. Use the guard to make sure tracing is enabled.
    Do this checking before we waste cycles gathering parameters to be traced.
    if (trcLogger.isLoggable(RASITraceEvent.TYPE_ENTRY_EXIT) {
        // since I want to trace 3 parameters, package them up in an Object[]
        Object[] parms = {new Integer(x), y, z};
        trcLogger.entry(RASITraceEvent.TYPE_ENTRY_EXIT, this, "methodX", parms);
    }
    ... logic
    // a debug or verbose trace point
    if (trcLogger.isLoggable(RASITraceEvent.TYPE_MISC_DATA) {
        trcLogger.trace(RASITraceEvent.TYPE_MISC_DATA, this, "methodX" "reached here");
    }
    ...
    // Another classification of trace event. Here an important state change
    has been detected, so a different trace type is used.
    if (trcLogger.isLoggable(RASITraceEvent.TYPE_SVC) {
        trcLogger.trace(RASITraceEvent.TYPE_SVC, this, "methodX", "an important event");
    }
    ...
    // ready to exit method, trace. No return value to trace
```

```

        if (trcLogger.isLoggable(RASITraceEvent.TYPE_ENTRY_EXIT)) {
            trcLogger.exit(RASITraceEvent.TYPE_ENTRY_EXIT, this, "methodX");
        }
    }
}

```

Setting up for integrated JRas operation

In the integrated mode of operation, message and trace events are sent to WebSphere Application Server logs. This is the default mode of operation.

1. Import the requisite JRas extensions classes

```

import com.ibm.ras.*;
import com.ibm.websphere.ras.*;

```

2. Declare logger references.

```

private RASMessageLogger msgLogger = null;
private RASTraceLogger trcLogger = null;

```

3. Obtain a reference to the Manager and create the loggers. Since loggers are named singletons, you can do this in a variety of places. One logical candidate for enterprise beans is the `ejbCreate()` method. For example, for the enterprise bean named "myTestBean", place the following code in the `ejbCreate()` method.

```

com.ibm.websphere.ras.Manager mgr = com.ibm.websphere.ras.Manager.getManager();
msgLogger = mgr.createRASMessageLogger("Acme", "WidgetCounter", "RasTest",
    myTestBean.class.getName());
// Configure the message logger to use the message file created
// for this application.
msgLogger.setMessageFile("acme.widgets.DefaultMessages");
trcLogger = mgr.createRASTraceLogger("Acme", "Widgets", "RasTest",
    myTestBean.class.getName());
mgr.addLoggerToGroup(trcLogger, groupName);

```

Setting up for combined JRas operation

In combined mode, messages and trace are logged to both WebSphere Application Server logs and user-defined logs. The following sample assumes that you have written a user defined handler named `SimpleFileHandler` and a user defined formatter named `SimpleFormatter`. It also assumes that you are not using user defined types or events.

1. Import the requisite JRas extensions classes

```

import com.ibm.ras.*;
import com.ibm.websphere.ras.*;

```

2. Import the user handler and formatter.

```

import com.ibm.ws.ras.test.user.*;

```

3. Declare the logger references.

```

private RASMessageLogger msgLogger = null;
private RASTraceLogger trcLogger = null;

```

4. Obtain a reference to the Manager, create the loggers and add the user handlers. Since loggers are named singletons, you can obtain a reference to the loggers in a number of places. One logical candidate for enterprise beans is the `ejbCreate()` method. Make sure that multiple instances of the same user handler are not accidentally inserted into the same logger. Your initialization code must handle this. The following sample is a message logger sample. The procedure for a trace logger is similar.

```

com.ibm.websphere.ras.Manager mgr = com.ibm.websphere.ras.Manager.getManager();
msgLogger = mgr.createRASMessageLogger("Acme", "WidgetCounter", "RasTest",
    myTestBean.class.getName());
// Configure the message logger to use the message file defined
// in the ResourceBundle sample.

```



```

msgLogger.setMessageFile("acme.widgets.DefaultMessages");

// Create the user handler and formatter. Configure the formatter,
// then add it to the handler.
RASHandler handler = new SimpleFileHandler("myHandler", "FileName");
RASFormatter formatter = new SimpleFormatter("simple formatter");
formatter.addEventClass("com.ibm.ras.RASMessageEvent");
handler.addFormatter(formatter);

// Add the Handler to the logger. Add the logger to the list of the
//handlers listeners, then set the handlers
// mask, which will update the loggers composite mask appropriately.
// WARNING - there is an order dependency here that must be followed.
msgLogger.addHandler(handler);
handler.addMaskChangeListener(msgLogger);
handler.setMessageMask(RASMessageEvent.DEFAULT_MESSAGE_MASK);

```

Setting up for stand-alone JRas operation

In stand-alone mode, messages and traces are logged only to user-defined logs. The following sample assumes that you have a user-defined handler named `SimpleFileHandler` and a user-defined formatter named `SimpleFormatter`. It is also assumed that no user-defined types or events are being used.

1. Import the requisite JRas extensions classes

```

import com.ibm.ras.*;
import com.ibm.websphere.ras.*;

```

2. Import the user handler and formatter.

```

import com.ibm.ws.ras.test.user.*;

```

3. Declare the logger references.

```

private RASMessageLogger msgLogger = null;
private RASTraceLogger trcLogger = null;

```

4. Obtain a reference to the Manager, create the loggers and add the user handlers. Since loggers are named singletons, you can obtain a reference to the loggers in a number of places. One logical candidate for enterprise beans is the `ejbCreate()` method. Make sure that multiple instances of the same user handler are not accidentally inserted into the same logger. Your initialization code must handle this. The following sample is a message logger sample. The procedure for a trace logger is similar.

```

com.ibm.websphere.ras.Manager mgr = com.ibm.websphere.ras.Manager.getManager();
msgLogger = mgr.createRASMessageLogger("Acme", "WidgetCounter", "RasTest",
    myTestBean.class.getName());
// Configure the message logger to use the message file defined in
//the ResourceBundle sample.
msgLogger.setMessageFile("acme.widgets.DefaultMessages");

// Get a reference to the Handler and remove it from the logger.
RASHandler aHandler = null;
Enumeration enum = msgLogger.getHandlers();
while (enum.hasMoreElements()) {
    aHandler = (RASHandler)enum.nextElement();
    if (aHandler instanceof WsHandler)
        msgLogger.removeHandler(wsHandler);
}

// Create the user handler and formatter. Configure the formatter,
// then add it to the handler.
RASHandler handler = new SimpleFileHandler("myHandler", "FileName");
RASFormatter formatter = new SimpleFormatter("simple formatter");
formatter.addEventClass("com.ibm.ras.RASMessageEvent");
handler.addFormatter(formatter);

```

```
// Add the Handler to the logger. Add the logger to the list of the
// handlers listeners, then set the handlers
// mask, which will update the loggers composite mask appropriately.
// WARNING - there is an order dependency here that must be followed.
msgLogger.addHandler(handler);
handler.addMaskChangeListener(msgLogger);
handler.setMessageMask(RASIMessageEvent.DEFAULT_MESSAGE_MASK);
```

Logging messages and trace data for Java server applications

By using the WebSphere Application Server for z/OS support for logging application messages and trace data, you can improve the reliability, availability, and serviceability of any Java application that runs in a WebSphere Application Server for z/OS server. Through this support, your Java application's messages can appear on the MVS master console, in the error log stream, or in the component trace (CTRACE) data set for WebSphere Application Server for z/OS. Your application's trace entries can appear in the same CTRACE data set.

Determining where to issue the messages

You might want to issue messages to the MVS master console to report serious error conditions for mission-critical applications. Through the master console, an operator can receive and, if necessary, take action in response to a message that indicates the status of an application. In addition, by directing messages to the master console, you can trigger automation packages to take action for specific conditions or events related to your application's processing.

Any messages that your application issues to the console also appear in either the error log stream or the CTRACE data set for WebSphere Application Server for z/OS, depending on the message type. Logging the messages in these system resources can help you more easily diagnose errors related to your application's processing. Similarly, issuing requests to log trace data in the CTRACE data set is another method of recording error conditions or collecting application data for diagnostic purposes.

System performance when logging messages and trace data

You can select the amount and types of trace data to be collected, which provides you with the ability to either run your application with minimal tracing when performance is a priority, or run your application with detailed tracing when you need to recreate a problem and collect additional diagnostic information.

The error log stream, the CTRACE data set for WebSphere Application Server for z/OS, and the master console are primarily intended for monitoring or recording diagnostic data for system components and critical applications. Depending on your installation's configuration, directing application messages and data to these resources might have an adverse affect on system performance. For example, if you send application data to the CTRACE data set, trace entries in that data set might wrap more quickly, which means you might lose some critical diagnostic data because the system writes new entries over existing ones when wrapping occurs. Use this logging support judiciously.

Note: You can only use WebSphere Application Server for z/OS support for logging messages and trace data for Java applications, not for Java applets.

Issuing application messages to the MVS master console

With the WebSphere Application Server for z/OS reliability, availability, and servicability support for Java (JRAS), you can issue messages from your Java application to the MVS master console. You might want to issue messages to the master console to report serious error conditions for mission-critical applications, or to trigger automation packages.

The messages your application issues also appear in either the error log stream or the component trace (CTRACE) data set that WebSphere Application Server for z/OS uses.

Logging the messages is another method of recording error conditions or collecting application data for diagnostic purposes.

Using a message logger

WebSphere Application Server for z/OS provides code that creates and manages a message logger, which processes your application's messages. WebSphere Application Server for z/OS creates only one message logger for each unique organization, product, or component, so that you can more easily identify the messages recorded in the error log stream or CTRACE data set for a specific application. The message logger runs in the Java Virtual Machine (JVM) for the WebSphere Application Server for z/OS server in which your Java application will run.

To use a message logger, in your Java application:

1. Define the message logger.
2. Drive the method to instruct WebSphere Application Server for z/OS to create the message logger.
3. Code messages at appropriate points in your application.

Chapter 4. Best practices for maintaining the run-time environment

Use the following guidelines to make sure that WebSphere Application Server for z/OS is customized and maintained correctly, to support your installation's application workload. Checking these basic software and hardware requirements can help you avoid problems with the run-time environment.

- **Check that you have the necessary prerequisite software up and running.** Check that they have the proper authorizations and that the definitions are correct.
- **Check for messages that signal potential problems.** Look for warning and error messages in the following sources:
 - SYSLOG from other z/OS subsystems and products, such as TCP/IP (especially the DNS, if in use), RACF, and so on
 - WebSphere Application Server for z/OS error log
 - SYSPRINT of the WebSphere Application Server for z/OS
 - Component trace (CTRACE) output for the server
- **Check the ports used by WebSphere Application Server.** The ports that are used by WebSphere Application Server must not be reserved by any other z/OS component.
- **Ensure that z/OS has enough DASD space for SVC dumps.** You might have to adjust the amount of space, because it depends on the size of your applications, on the configured Java virtual machine (JVM) heap size, and on the number of servant regions that might be included in one dump, and so on. For an SVC dump of one controller and one servant, you can start with a minimum of 512, but might have to increase the MAXSPACE to 1024 or higher, given the factors listed above.
- **Check your general environment.** Does your system have enough memory? Insufficient memory problems can show up as AUX shortages, abends, or exceptions from the WebSphere Application Server for z/OS run-time. Sometimes the heap size for Language Environment (LE) and for the Java virtual machine (JVM) needs to be increased. If you are using RRS and DB2, make sure your system has enough space for archive data sets.
- **Make sure all prerequisite fixes have been installed;** a quick check for a fix can save hours of debugging.

For the most current information on fixes and service updates, see:

- The Preventive Service Planning (PSP) buckets for both WebSphere Application Server for z/OS and JAVA subsets of the WebSphere Application Server for z/OS Upgrade. To obtain a copy of the most current versions of these PSP buckets, you can either contact the IBM Support Center, use S/390 SoftwareXcel or link to IBMLink.
- The Support Web page of the WebSphere Application Server for z/OS Web site, which contains a table of the latest authorized program analysis reports (APARs).

With the latest service information, check the following:

- Ensure that all prerequisite PTFs (fixes) have been applied to the system.
- Verify that all PTFs were actually present in the executables that were used at the time of error. Often, SMP can indicate that a fix is present and installed on the system when, in reality, the executables that were used at the time of error did not contain the fix.

Best practices for using system controls

- You have the option of using a z/OS system logger log stream as the WebSphere Application Server for z/OS error log. The WebSphere variable `ras_log_logstreamName`

identifies which log stream you want to use for the error log; it has no default setting. If you do not use a log stream, however, messages that usually appear in the error log are directed to server's job log.

- You have the option of directing trace output to SYSPRINT or buffers. The WebSphere variable `ras_trace_outputLocation`

controls the location of trace output; its default values are SYSPRINT for client applications, and buffers to all other processes. Although you can change the default for other processes from buffers to SYSPRINT, performance is better when you use buffers.

- You can use the Resource Measurement Facility (RMF) to view status information that might indicate potential problems. WebSphere Application Server for z/OS uses Workload Manager (WLM) services to report transaction begin-to-end response times and execution delay times, which might indicate that changes are required for timeout values or tuning controls.

Updating the CFRM policy

You must update the coupling facility resource management (CFRM) policy before using log streams that are CF-resident, such as the WebSphere error log and RRS logs. If you have the source for the current active CFRM policy, update the source and use the IXCMIAPU Administrative Data utility to generate the new policy. If you do not have the source for the current active CFRM policy, rebuild the source from the active CFRM policy:

1. Find the active policy by issuing the command: `D XCF,POL` You will get output similar to this (partial display):

```
D XCF,POL
IXC364I 10.57.49 DISPLAY XCF 061

. . .
TYPE: CFRM
POLNAME: POLCF1N1
STARTED: 03/14/2003 11:32:22
LAST UPDATED: 03/14/2003 11:31:52

. . .
```

2. List the active CFRM Policy's structure definitions by using the Administrative Data utility:

```
//STEP1 EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=*
//SYSABEND DD SYSOUT=*
//SYSIN DD *
DATA TYPE(CFRM) REPORT(YES)
/*
```

3. Extract the definitions for the ACTIVE policy only. Using the SYSPRINT from the above utility job, edit the output with the following steps so it can be used to define a new policy in the next job:
 - Extract the definitions for the ACTIVE policy only.
 - Delete the heading lines.

- Add the new structure definition using the BBOWCFRM member of the target CNTL dataset as a model.
 - Copy it into a FB-LRECL(80) dataset to be used as SYSIN for the following job.
4. Use the Administrative Data utility to update the CFRM policy. The policy name can be the same as the ACTIVE CFRM policy or a new name. If you use the active policy name, REPLACE(YES) must be specified on the DEFINE control statement.

```
//STEP20 EXEC PGM=IXCMIAPU
//SYSPRINT DD SYSOUT=*
//SYSABEND DD SYSOUT=*
//SYSIN DD *
DATA TYPE(CFRM) REPORT(YES)

DEFINE POLICY NAME(POLCF1N1) REPLACE(YES)

CF NAME(CF1LPAR) DUMPSPACE(5000) PARTITION(0E) CPCID(00)
TYPE(009672) MFG(IBM) PLANT(02) SEQUENCE(000000051205)

CF NAME(CF2LPAR) DUMPSPACE(5000) PARTITION(0F) CPCID(00)
TYPE(009672) MFG(IBM) PLANT(02) SEQUENCE(000000051205)

STRUCTURE NAME(CTS130_DFHLOG) SIZE(24000) INITSIZE(12000)
REBUILDPERCENT(1) PREFLIST(CF1LPAR, CF2LPAR)
. . .
<== Insert your new structure definitions here
```

5. Activate the new policy by issuing the following MVS Command:
- ```
SETXCF START,POLICY,TYPE=CFRM,POLNAME=POLCF1N1
```

---

## Configuring WebSphere Application Server for z/OS variables

WebSphere Application Server for z/OS provides configuration variables that control server behavior. Specifically, these variables allow you to control:

- Output destinations and characteristics for the error log, and for CTRACE buffers, data sets and the external writer.
- Trace buffers, data sets, and the content of trace data.
- Types of dumps to be requested.
- Timeout values for system and application behavior.

Generally speaking, the default values are designed for normal operation in a production environment. Other circumstances might require different values:

- When you first customize and verify WebSphere Application Server for z/OS installation, or
- When you test application workloads in a test environment, or
- when you encounter a problem, and need to collect more diagnostic data.

### Steps for configuring WebSphere variables

You should know that:

- Configuration variables may be set on a cell, node, or server level.
  - Variable values set on a cell level apply to all servers in all nodes in the cell, unless a different value for the same variable is set on a node or server level. Variable settings on a node or server level override values for the same variable set at the cell level.

- Variables set on a node level apply to all servers in the node, unless a different value for the same variable is set on the server level. Variable settings on a server level override values for the same variable set at the node or cell level.
- Variables set on a server level apply only to the specific server, not to any other servers in the same node or cell.

When you are diagnosing particular problems, you are most likely to alter variable values on a server level, for a particular server. Specifying variable values on the server level affects both the controller and servant regions.

- You may use scripting interfaces, instead of the WebSphere Administrative console, to alter configuration variable values.

Depending on the types of problems you encounter, you might need to change the values set for WebSphere configuration variables that control server behavior. The following procedure explains how to use the WebSphere Administrative console to change configuration variable values, commonly known as console settings.

1. Click **Environment -> Manage WebSphere Variables** in the console navigation tree.
2. On the **WebSphere Variables** page, select **Server** as the scope of the variable setting, and click **Apply**.
3. On the **WebSphere Variables** page, click **New**.
4. On the **Variable** page, specify a name and value for the variable. So other people can understand what the variable is used for, also specify a description for the variable. Then click **OK**.
5. Verify that the variable is shown in the list of variables.
6. Save your configuration.
7. To have the configuration take effect, stop the server and then start the server again.

---

## Setting up component trace (CTRACE)

This file describes how to set up CTRACE

WebSphere Application Server for z/OS uses z/OS component trace (CTRACE) facilities to manage the collection and storage of trace data. Unless you configure specific CTRACE controls, WebSphere Application Server for z/OS records its trace data in address-space buffers in private (pageable) storage. This data is not accessible unless a dump of the address space is taken.

Although CTRACE data is primarily output for IBM service personnel to use, exploiting CTRACE capabilities at your installation allows you to have additional trace data available when a problem first occurs. Because CTRACE efficiently uses system resources, you can collect valuable trace data with minimal impact on performance.

When your installation first customizes and verifies WebSphere Application Server for z/OS installation, you have the option of defining CTRACE controls and resources. Using the ISPF customization dialog to configure a base application server node, you can specify:

- Data sets to contain CTRACE data collected for WebSphere Application Server for z/OS.
- CTRACE writer parameters that control the writer through which trace data moves from address-space buffers into trace data sets.



- The parmlib member that connects WebSphere Application Server for z/OS address spaces to trace data sets, and optionally turns on the CTRACE writer.
- WebSphere variables that control the characteristics of trace data.

The ISPF customization dialog generates instructions for:

1. Starting the CTRACE writer
2. Starting the WebSphere Application Server for z/OS application server

Following the instructions in sequence is quite important; you can lose valuable trace data if you do not start the CTRACE writer before starting the server.

For information about the advantages of using CTRACE facilities, see *z/OS MVS Diagnosis: Tools and Service Aids, GA22-7589*

## Steps for preparing CTRACE controls and resources

Before you start CTRACE activity for WebSphere Application Server for z/OS servers, you need to make some decisions about CTRACE controls and resources. You have the option of using default CTRACE values and resources, such as the IBM-supplied CTRACE parmlib member for WebSphere Application Server for z/OS, or you may alter default values and provide your own resources.

Perform the following steps to prepare CTRACE controls and resources:

1. Decide where you want to store CTRACE data. The location of trace data is set through the WebSphere variable `ras_trace_outputLocation`. Make a note of the value that you want to use; you will set the value later, when you start CTRACE activity.

### Note: Tips:

- See *Setting Trace Controls* for an explanation of and default values for the `ras_trace_outputLocation` variable.
  - If you decide to use trace data sets, you can use existing or create new data sets now or later, as part of the WebSphere Application Server for z/OS customization process.
  - If you want the CTRACE data for each cell to go into separate data sets, use the `ras_trace_ctraceParms` variable described in *Setting Trace Controls*.
  - When you are installing WebSphere Application Server for z/OS, sending trace data to `SYSPRINT` can be helpful; however, tracing to `SYSPRINT` in a production environment can cause spool space to fill up quickly. For production, you can specify a different trace output location through the WebSphere variable `ras_trace_outputLocation`.
  - To separate trace data from other standard output messages, such as `System.output.println` from Java applications, you can direct CTRACE data to a separate data set. To accomplish this separation, you need to:
    - Specify a `TRCFILE` DD statement in the JCL start procedure (`proc`) for the server.
    - Set the WebSphere variable `ras_trace_outputLocation` to `TRCFILE`. Note that you may specify the `TRCFILE` value with or without additional variable values.
2. Through modify commands, you have the ability to dynamically and temporarily direct trace data to `SYSPRINT` or `TRCFILE`.

- a. You can direct CTRACE data to buffers as part of normal operations.
- b. When an error occurs, you can use the modify command to send trace data to SYSPRINT or TRCFILE.
- c. Then you can use the modify command again, to dynamically reset the trace-output location

You can complete this sequence of actions without stopping the server that is encountering the problem.

3. Decide whether you want to accept defaults or provide other values for the following WebSphere variables. Make a note of the values that you want to use; you will set the values later, when you start CTRACE activity.
  - ras\_time\_local
  - ras\_trace\_BufferCount
  - ras\_trace\_BufferSize

For defaults and valid values, see Setting Trace Controls

4. Decide whether you want to use the default JCL start procedure for the CTRACE writer, or code your own JCL proc. WebSphere Application Server for z/OS provides a CTRACE writer proc named bbowtr. If you decide to provide your own CTRACE writer procedure, create the JCL start proc using the rules for source JCL for an external writer in . Place the start procedure in your system proclib.
5. Decide whether you want to use the default CTRACE parmlib member for WebSphere Application Server for z/OS, or provide your own. The WebSphere parmlib member is named CTIBBO00. If you decide to provide your own parmlib member, create one according to the following rules, and place it in your system parmlib. **Rules:**
  - You must follow the same naming convention; that is, the name must consist of the letters CTIBB0, but you may replace the two zeroes with any two alphanumeric characters.
  - Your parmlib member must contain the following parameters:

```
TRACEOPTS WTRSTART(xxxxxx)
ON /*CONNECT TO CTRACE EXTERNAL WRITER: */ WTR(xxxxxx)
```

Table 1. Parameters for the CTIBBOxx parmlib member

| Parameter         | Required or Optional?   | Description                                                                                                                                                                                                                                                                                                                                                                                     |
|-------------------|-------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| TRACEOPTS         | Required and positional | This parameter must be specified first.                                                                                                                                                                                                                                                                                                                                                         |
| WTRSTART (xxxxxx) | Optional                | This parameter automatically starts the CTRACE writer, using the specified JCL procedure.<br><br><b>Alternative:</b> You may use the operator command TRACE CT,WTRSTART=xxxxx to start the CTRACE writer. If you use this parameter in the parmlib member, and later issue the operator command, CTRACE issues an informational message to notify you that the writer already has been started. |
| ON                | Required                | This parameter must be specified before component options.                                                                                                                                                                                                                                                                                                                                      |

Table 1. Parameters for the CTIBBOxx parmlib member (continued)

|             |          |                                                                                                                                                                                                                                                       |
|-------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| WTR(XXXXXX) | Required | This parameter identifies the JCL procedure to be used to start the CTRACE external writer. The specified value must match the value of the WTRSTART parameter to capture the WebSphere Application Server for z/OS trace data into a trace data set. |
|-------------|----------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

6. Ensure that the DLL named BBORTSS5 is in the link-pack area (LPA).

After you have made these decisions and completed preparations, you are ready to start CTRACE activities using one of the following procedures:

- During customization of WebSphere Application Server for z/OS (see “Steps for starting CTRACE as part of WebSphere Application Server for z/OS customization” ) or
- While WebSphere Application Server for z/OS servers already are running on z/OS or OS/390 (see “Steps for starting CTRACE while WebSphere Application Server for z/OS servers are active” on page 38

## Steps for starting CTRACE as part of WebSphere Application Server for z/OS customization

Make sure you have properly prepared CTRACE controls and resources as described in “Steps for preparing CTRACE controls and resources” on page 35

Perform the following steps to start CTRACE as part of the customization process for WebSphere Application Server for z/OS:

1. Using the ISPF customization dialog to configure a base application server node, specify: Trace data set characteristics, CTRACE writer parameters, the CTRACE parmlib member and WebSphere variables, if you want values other than the defaults. The ISPF customization dialog generates instructions for starting the CTRACE writer and instructions for starting the WebSphere Application Server for z/OS application server
2. Start the CTRACE writer, using the generated instructions. You must start the writer first, or you might lose valuable trace data.
3. Start the WebSphere Application Server for z/OS application server, using the generated instructions.
4. When you need to collect trace data for analysis:
  - a. Use the following operator command to disconnect WebSphere Application Server for z/OS from CTRACE: TRACE CT,ON,COMP= *cell\_short\_name* REPLY x,WTR=DISCONNECT,END. where *cell\_short\_name* is the value specified through the ISPF customization dialog to identify the location of server configuration files. The name must be 8 or fewer characters and all uppercase.
  - b. Use the following operator command to stop the CTRACE writer address space: TRACE CT,WTRSTOP= *writer\_name* where *writer\_name* is the name of the JCL start procedure for the CTRACE writer.

## Steps for starting CTRACE while WebSphere Application Server for z/OS servers are active

Make sure you have properly prepared CTRACE controls and resources as described in “Steps for preparing CTRACE controls and resources” on page 35

If you start a WebSphere Application Server for z/OS server before starting the CTRACE writer for WebSphere, the server still gathers data in its trace buffers. This trace data is not available for use unless you follow this procedure, or until a dump of the server address space is taken.

Perform the following steps to start CTRACE when a WebSphere Application Server for z/OS server already is active:

1.
  - Use the following operator command to start the CTRACE writer address space: TRACE CT,WTRSTART= *writer\_name* where *writer\_name* is the name of the JCL start procedure for the CTRACE writer that is specified in the WebSphere Application Server for z/OS CTIBBO *xx* parmlib member.
  - To connect WebSphere Application Server for z/OS to a CTRACE writer other than the one specified in the CTIBBO *xx* parmlib member, also enter these operator commands: TRACE CT,ON,COMP= *cell\_short\_name* REPLY x,WTR= *writer\_name*, END where *cell\_short\_name* is the value specified through the ISPF customization dialog to identify the location of server configuration files (the name must be 8 or fewer characters and all uppercase) and *writer\_name* is the name of a JCL start procedure for a CTRACE external writer. The JCL start procedure must reside in the system’s proclib.

The CTRACE external writer begins writing the server’s trace data to the location specified through the WebSphere variable *ras\_trace\_outputLocation*

2. When you need to collect trace data for analysis:
  - a. Use the following operator command to disconnect WebSphere Application Server for z/OS from CTRACE: TRACE CT,ON,COMP= *cell\_short\_name* REPLY x,WTR=DISCONNECT,END
  - b. Use the following operator command to stop the CTRACE writer address space: TRACE CT,WTRSTOP= *writer\_name* where *writer\_name* is the name of the JCL start procedure for the CTRACE writer.

## Using CTRACE to collect trace data for Java server applications

This file describes how to set up CTRACE

Applications that run in WebSphere Application Server for z/OS can use JRas to provide tracing support that is consistent with WebSphere tracing. Instrumented applications use the JRas interfaces and classes for logging and tracing; trace data is written to the same component trace data set as the internal traces issued by the WebSphere Application Server for z/OS runtime. So you can gather application trace data in the same locations, and use the same commands to start and stop CTRACE for these JRas applications as you do for WebSphere Application Server for z/OS server in which the applications are running.

## Setting up the error log

WebSphere Application Server for z/OS uses an error log to record error information when an unexpected condition or failure is detected within the product's own code. Such unexpected conditions or failures include:

- Assertion failures
- Unrecoverable error conditions
- Failures related to vital resources, such as memory
- Operating system exceptions
- Programming defects in WebSphere Application Server for z/OS code.

Because WebSphere Application Server for z/OS is predefined as a z/OS system logger application, you can use a log stream as the product's error log. Doing so offers the following flexibility:

- You can direct error information to:
  - A coupling facility log stream, which provides sysplex-wide error logging, or
  - A DASD-only log stream, which provides single system-only error logging.
- You can set up a common log stream for all WebSphere Application Server for z/OS servers, or individual log streams for each application server. Local z/OS or OS/390 client ORBs can also log data in log streams. The system logger APIs are unauthorized, but log stream resources can be protected using security products such as RACF.
- You can use the WebSphere variable  
`ras_time_local`

to control whether timestamps in the error log appear in local time or Greenwich Mean Time (GMT), which is the default.

When your installation first customizes and verifies WebSphere Application Server for z/OS installation, you have the option of defining the error log as a log stream. Using the ISPF customization dialog to configure a base application server node, you can specify log stream characteristics, including sizes. After verifying installation, you can change the log stream used for normal operations.

For additional information about z/OS log stream requirements, access the z/OS *MVS Setting up a Sysplex*, SA22-7625 available on the z/OS Library Web page

## Using the z/OS modify command

This file lists parameters that allow you to dynamically change values related to tracing activity for a server or servant.

You may use either the WebSphere Application Server administrative console or the z/OS MVS console to accomplish many operator tasks related to WebSphere Application Server for z/OS servers. Entering the z/OS display or modify commands through the MVS console can provide information or perform tasks that are useful for diagnosing problems.

In particular, the z/OS modify command provides parameters that not only allow you to control WebSphere Application Server for z/OS operations, but also to:

- Display information about WebSphere Application Server for z/OS servers or servants (regions), and
- Dynamically change values related to tracing activity for a server or servant.

Recent additions to the modify command enable you to display information about the following:

- Active controllers (servers)
- Servants
- Sessions
- Trace settings
- Java trace
- JVM heap statistics
- Work Elements
- Error Log

### Displaying the options on the modify command

You can display the options on the modify command with the help parameter. For example, the command

```
f azsr01a,help
```

generates the following output:

```
BB000178I THE COMMAND MODIFY MAY BE FOLLOWED BY ONE OF THE FOLLOWING KEYWORDS:
BB000179I CANCEL - CANCEL THIS CONTROL REGION
BB000179I TRACEALL - SET OVERALL TRACE LEVEL
BB000179I TRACEBASIC - SET BASIC TRACE COMPONENTS
BB000179I TRACEDetail - SET DETAILED TRACE COMPONENTS
BB000179I TRACESPECIFIC - SET SPECIFIC TRACE POINTS
BB000179I TRACEINIT - RESET TO INITIAL TRACE SETTINGS
BB000179I TRACENONE - TURN OFF ALL TRACING
BB000179I TRACETOSYSPRINT - SEND TRACE OUTPUT TO SYSPRINT (YES/NO)
BB000179I TRACETOTRCFILE - SEND TRACE OUTPUT TO TRCFILE DD CARD (YES/NO)
BB000179I DISPLAY - DISPLAY STATUS
BB000179I TRACE_EXCLUDE_SPECIFIC - EXCLUDE SPECIFIC TRACE POINTS
BB000179I TRACEJAVA - SET JAVA TRACE OPTIONS
```

You can display the DISPLAY options with the display,help parameter. For example, the command

```
f azsr01a,display,help
```

generates the following output:

```
BB000178I THE COMMAND DISPLAY, MAY BE FOLLOWED BY ONE OF THE FOLLOWING KEYWORDS:
BB000179I SERVERS - DISPLAY ACTIVE CONTROL PROCESSES
BB000179I SERVANTS - DISPLAY SERVANT PROCESSES OWNED BY THIS CONTROL PROCESS
BB000179I SESSIONS - DISPLAY INFORMATION ABOUT COMMUNICATIONS SESSIONS
BB000179I TRACE - DISPLAY INFORMATION ABOUT TRACE SETTINGS
BB000179I JVMHEAP - DISPLAY JVM HEAP STATISTICS
BB000179I WORK - DISPLAY WORK ELEMENTS
BB000179I ERRLOG - DISPLAY THE LAST 10 ENTRIES IN THE ERROR LOG
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,HELP
```

### Displaying work elements

The z/OS `modifyserver,display,work` command can provide useful information for diagnosing problems, or displaying what is going on inside your application servers. The syntax is:

```
f servername,display,work,display work parameters
```

where the **display work parameters** can be one of the following:

**HELP**

displays the DISPLAY,WORK parameters

**EJB**

EJB requests driven by IIOP: total, current, dispatched & timed out

**EJB,SRS**

EJB requests driven by IIOP by servant region

**SERVLET**

Servlet requests driven by HTTP: total, current, dispatched & timed out

**SERVLET,SRS**

Servlet requests by each servant region

**MDB**

MDB requests driven by JMS: total, current, dispatched & timed out

**MDB,SRS**

MDB requests broken down by servant region

**ALL**

Combines the above for EJBs, servlets, and MDBs

**ALL,SRS**

EJBs, servants, and MDBs by servlets region

**SUMMARY**

total requests, current in progress, and in dispatch for all types

**SUMMARY,SRS**

total received by each servant region, current in dispatch in each servant region for all types

The following examples demonstrate how to use various display,work parameters and the resulting output:

```
f azsr01a,display,work,servlet
```

```
BB000255I TIME OF LAST WORK DISPLAY Wed Dec 3 19:17:54 2003
BB000256I TOTAL SERVLET REQUESTS 150670 (DELTA 1654)
BB000257I CURRENT SERVLET REQUESTS 1
BB000258I SERVLET REQUESTS IN DISPATCH 0
BB000267I TOTAL SERVLET TIMEOUTS 0 (DELTA 0)
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,WORK,SERVLET
```

```
f azsr01a,display,work,servlet,srs
```

```
BB000255I TIME OF LAST WORK DISPLAY Wed Dec 3 19:18:01 2003
BB000259I STC18964: TOTAL SERVLET REQUESTS 152344 (DELTA 1675)
BB000260I STC18964: SERVLET REQUESTS IN DISPATCH 0
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,WORK,SERVLET,ALL
```

(EJB and MDB displays would look the same except for the request type.)

```
f azsr01a,display,work,summary
```

```
BB000255I TIME OF LAST WORK DISPLAY Wed Dec 3 19:18:38 2003
BB000261I TOTAL REQUESTS TO SERVER 173591 (DELTA 13944)
BB000262I TOTAL CURRENT REQUESTS 0
BB000263I TOTAL REQUESTS IN DISPATCH 0
BB000268I TOTAL TIMED OUT REQUESTS 0 (DELTA 0)
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,WORK,SUMMARY
```

```
f azsr01a,display,work,summary,srs
```

```
BB000255I TIME OF LAST WORK DISPLAY Wed Dec 3 19:27:01 2003
```



```

BB000264I STC18964: TOTAL REQUESTS 173591 (DELTA 0)
BB000265I STC18964: TOTAL REQUESTS IN DISPATCH 0
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,WORK,SUMMARY,SRS

```

To display a combination of all information with one command, use the all keyword, as follows:

```
f azsr01a,display,work,all
```

```

BB000255I TIME OF LAST WORK DISPLAY Wed Dec 3 19:27:07 2003
BB000256I TOTAL EJB REQUESTS 0 (DELTA 0)
BB000257I CURRENT EJB REQUESTS 0
BB000258I EJB REQUESTS IN DISPATCH 0
BB000267I TOTAL EJB TIMEOUTS 0 (DELTA 0)
BB000256I TOTAL SERVLET REQUESTS 173591 (DELTA 0)
BB000257I CURRENT SERVLET REQUESTS 0
BB000258I SERVLET REQUESTS IN DISPATCH 0
BB000267I TOTAL SERVLET TIMEOUTS 0 (DELTA 0)
BB000256I TOTAL MDB REQUESTS 0 (DELTA 0)
BB000257I CURRENT MDB REQUESTS 0
BB000258I MDB REQUESTS IN DISPATCH 0
BB000267I TOTAL MDB TIMEOUTS 0 (DELTA 0)
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,WORK,ALL

```

```
f azsr01a,display,work,all,srs
```

```

BB000255I TIME OF LAST WORK DISPLAY Wed Dec 3 19:30:06 2003
BB000259I STC18964: TOTAL EJB REQUESTS 0 (DELTA 0)
BB000260I STC18964: EJB REQUESTS IN DISPATCH 0
BB000259I STC18964: TOTAL SERVLET REQUESTS 173591 (DELTA 0)
BB000260I STC18964: SERVLET REQUESTS IN DISPATCH 0
BB000259I STC18964: TOTAL MDB REQUESTS 0 (DELTA 0)
BB000260I STC18964: MDB REQUESTS IN DISPATCH 0
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,WORK,ALL,SRS

```

**Note:** Some points to consider:

- Adding,HELP at the end will tell you what your choices are (for example, DISPLAY,WORK,HELP).
- DISPLAY,WORK produces the same output as DISPLAY,WORK,SUMMARY.
- The Display EJB command shows you IOP requests. So, if you drive an HTTP request that runs a JSP or servlet that calls an EJB, only the servlet count increments. The counts are by protocol used to get to the CTL region.
- When displaying the work by servant region, only total and "in dispatch" requests are shown. Without the SRS keyword you will also see the current requests and time outs.
- Everything is not frozen while the numbers are collected. The "old" data is saved away and used to calculate the "delta" in the next command. (Under a heavy load, the displayed values and deltas may not add up as you would expect.)
- TOTAL requests will also show the delta since last DISPLAY,WORK command, but it will not include the current "in flight" requests.

### Displaying the error log

You can display the contents of the product's error log with the DISPLAY,ERRLOG option of the modify command. This shows the last ten messages in the error log even if you are not routing them to a log stream. For example, the command

```
f x5sr01b,display,errlog
```



generates the following output:

```

BB000266I (STC18876) BossLog: { 0001} 2003/11/25 20:08:55.120 01
SYSTEM=SYSB SERVER=X5SR01B PID=0X010201B2 TID=0X12FB3F00 00000000
c=UNK ./bborjtr.cpp+812 ... BB000222I TRAS0017I: The startup trace
state is *=all=disabled.
BB000266I (STC18876) BossLog: { 0002} 2003/11/25 20:09:08.255 01
SYSTEM=SYSB SERVER=X5SR01B PID=0X010201B2 TID=0X12FB3F00 00000000
c=UNK ./bborjtr.cpp+812 ... BB000222I SECJ0231I: The Security
component's FFDC Diagnostic Module com.ibm.ws.security.core.SecurityDM
registered successfully: true.
BB000266I (STC18876) BossLog: { 0003} 2003/11/25 20:09:09.562 01
SYSTEM=SYSB SERVER=X5SR01B PID=0X010201B2 TID=0X12FB3F00 00000000
c=UNK ./bborjtr.cpp+812 ... BB000222I SECJ0212I: WCCM JAAS
configuration information successfully pushed to login provider class.
BB000266I (STC18876) BossLog: { 0004} 2003/11/25 20:09:09.573 01
SYSTEM=SYSB SERVER=X5SR01B PID=0X010201B2 TID=0X12FB3F00 00000000
c=UNK ./bborjtr.cpp+812 ... BB000222I SECJ0240I: Security service
initialization completed successfully
BB000266I (STC18876) BossLog: { 0005} 2003/11/25 20:09:18.304 01
SYSTEM=SYSB SERVER=X5SR01B PID=0X010201B2 TID=0X12FB3F00 00000000

c=UNK ./bborjtr.cpp+812 ... BB000223I PMI0023W: Unable to register
PMI module due to duplicate name: SoapConnectorThreadPool
BB000266I (STC18876) BossLog: { 0006} 2003/11/25 20:09:29.451 01
SYSTEM=SYSB SERVER=X5SR01B PID=0X010201B2 TID=0X12FB3F00 00000000
c=UNK ./bborjtr.cpp+812 ... BB000222I SECJ0243I: Security service
started successfully
BB000266I (STC18876) BossLog: { 0007} 2003/11/25 20:09:29.464 01
SYSTEM=SYSB SERVER=X5SR01B PID=0X010201B2 TID=0X12FB3F00 00000000
c=UNK ./bborjtr.cpp+812 ... BB000222I SECJ0210I: Security enabled false
BB000266I (STC18876) BossLog: { 0008} 2003/11/25 20:09:35.772 01
SYSTEM=SYSB SERVER=X5SR01B PID=0X010201B2 TID=0X12FB3F00 00000000
c=UNK ./bborjtr.cpp+812 ... BB000223I PMI0023W: Unable to register PMI
module due to duplicate name: ProcessDiscovery
.
BB000188I END OF OUTPUT FOR COMMAND DISPLAY,ERRLOG

```

## z/OS modify command parameters and their equivalent WebSphere variables

The z/OS modify command provides parameters that allow you to dynamically change values related to tracing activity for a server or servant. Table 2 lists the modify command parameters and the WebSphere variable that provides equivalent function.

Table 2. z/OS modify command parameters and their equivalent WebSphere variables

| z/OS modify command parameter | Equivalent WebSphere variable                                                                  | For more information, see..                  |
|-------------------------------|------------------------------------------------------------------------------------------------|----------------------------------------------|
| TRACEALL                      | ras_trace_defaultTracingLevel                                                                  | Internal tracing tips for WebSphere for z/OS |
| TRACEBASIC                    | ras_trace_basic Note: Do not change this variable unless directed by IBM service personnel.    | Setting trace controls for IBM service       |
| TRACEDetail                   | ras_trace_detail                                                                               | Setting trace controls for IBM service       |
| TRACESPECIFIC                 | ras_trace_specific Note: Do not change this variable unless directed by IBM service personnel. | Setting trace controls for IBM service       |

Table 2. z/OS modify command parameters and their equivalent WebSphere variables (continued)

|                        |                                                                                                           |                                              |
|------------------------|-----------------------------------------------------------------------------------------------------------|----------------------------------------------|
| TRACE_EXCLUDE_SPECIFIC | rac_trace_exclude_specific<br>Note: Do not change this variable unless directed by IBM service personnel. | Setting trace controls for IBM service       |
| TRACEINIT              | n/a                                                                                                       | Example: Getting help for the modify command |
| TRACENONE              | n/a                                                                                                       | Example: Getting help for the modify command |
| TRACETOSYSPPRINT       | ras_trace_outputLocation=SYSPPRINT                                                                        | Internal tracing tips for WebSphere for z/OS |
| TRACETOTRCFILE         | ras_trace_outputLocation=TRCFIL                                                                           | Internal tracing tips for WebSphere for z/OS |
| TRACEJAVA              | n/a                                                                                                       | Example: Getting help for the modify command |

## Types of configuration variables

WebSphere Application Server for z/OS provides configuration variables that allow you to control:

- Output destinations and characteristics for the error log, and for CTRACE buffers, data sets and the external writer.
- Trace buffers, data sets, and the content of trace data.
- Types of dumps to be requested.
- Timeout values for system and application behavior.

## Setting output destinations and characteristics

**client\_ras\_logstreamname= *name*** Specifies the name of the log stream for an application client run-time to use for error information.

**Default:** If this variable is not specified, the client run-time uses STDERR.

**Example:**

```
client_ras_logstreamname=MY.CLIENT.ERROR.LOG
```

**Tip:** Do not put the log stream name in quotes. Log stream names are not data set names.

**ras\_default\_msg\_dd= *DD\_card\_name*** Redirects write-to-operator (WTO) messages that use the default routing to hardcopy. These messages are redirected to the location identified through the DD card on the server's JCL start procedure. These WTO messages are primarily messages that WebSphere Application Server for z/OS issues during initialization.

**Default:** No default value is set; WTO messages that use default routing are sent to hardcopy.

**Examples:**

```
ras_default_msg_dd=MSGDD
ras_default_msg_dd=DFLTLOG
```

Example of the

DFLTLOG DD

card on the server's JCL start procedure:

```
//DFLTLOG DD SYSOUT=*
```

**ras\_hardcopy\_msg\_dd= *DD\_card\_name*** Redirects write-to-operator (WTO) messages that WebSphere Application Server for z/OS routes to hardcopy. These messages are redirected to the location identified through the DD card on the server's JCL start procedure. These WTO messages are primarily audit messages issued from Java code during initialization.

**Default:** No default value is set; WTO messages that use hardcopy routing are sent to hardcopy.

**Example:**

```
ras_hardcopy_msg_dd=MSGDD
```

**ras\_log\_logstreamName= *name*** Specifies the log stream for WebSphere Application Server for z/OS to use for error information during bootstrap processing. If the specified log stream is not found or not accessible, a message is issued and errors are written to the server's job log.

**Default:** If this variable is not specified, WebSphere Application Server for z/OS uses STDERR.

**Example:**

```
ras_log_logstreamName=MY.CB.ERROR.LOG
```

**Tip:** Do not put the log stream name in quotes. Log stream names are not data set names.

## Setting trace controls

Controlling behavior through timeout values

**ras\_trace\_outputLocation=SYSPRINT | BUFFER | TRCFILE**

Specifies where you want trace records to be sent:

- To SYSPRINT
- To a memory buffer (BUFFER), the contents of which are later written to a CTRACE data set
- To a trace data set (TRCFILE) specified on the TRCFILE DD statement in the server's start procedure.

For servers, you may specify one or more values, separated by a space. For clients, you may specify SYSPRINT only.

**Defaults:**

- For clients, SYSPRINT
- For all other processes, BUFFER

**Example:** ras\_trace\_outputLocation=SYSPRINT BUFFER

**ras\_time\_local=0 | 1**

Specifies whether timestamps in trace records use Greenwich Mean Time (GMT) or local time. This variable setting controls timestamp format in the error log, and in traces sent to SYSPRINT or TRCFILE DD.

**Default:** 0 (GMT)

**Example:** `ras_time_local=1` sets timestamps to local time.

**ras\_trace\_ctraceParms=SUFFIX | MEMBER\_NAME**

Identifies the CTRACE PARMLIB member. The value can be either:

- A two-character suffix, which is added to the string CTIBB0 to form the name of the PARMLIB member, or
- The fully specified name of the PARMLIB member. A fully specified name must conform to the naming requirements for a CTRACE PARMLIB member.

If this environment variable is specified and the PARMLIB member is not found, the default PARMLIB member, CTIBBO00, is used. If neither the specified nor the default PARMLIB member is found, tracing is defined to CTRACE, but there is no connection to a CTRACE external writer.

**Note:** The Daemon is the only server that recognizes this environment variable.

**Default:** 00 (the default PARMLIB member, CTIBBO00)

**Example:** `ras_trace_ctraceParms=01` identifies PARMLIB member CTIBBO01

**ras\_trace\_BufferCount= *n***

Specifies the number of trace buffers to allocate. Valid values are 4 through 8.

**Default:** 4

**Example:** `ras_trace_BufferCount=6`

**ras\_trace\_BufferSize= *n***

Specifies the size of a single trace buffer in bytes. You can use the letters K (for kilobytes) or M (for megabytes). Valid values are 128K through 4M.

**Default:** 1M

**Example:** `ras_trace_BufferSize=2M`

## Setting dump controls

**ras\_dumpoptions\_dumptype= *n***

Specifies the default dump used by the signal handler. Valid values and their meanings are:

- 0  
No dump is generated.
- 1  
A ctrace dump is taken.
- 2  
A cdump dump is taken.
- 3  
A csnap dump is taken.
- 4  
A CEE3DMP dump is taken.

CEE3DMP generates a dump of Language Environment and the member language libraries. Sections of the dump are selectively included, depending on dump options specified, either by default or through the

`ras_dumpoptions_ledumpoptions`

variable. By default, this value passes

THREAD(ALL) BLOCKS

to CEE3DMP. You can override the default options for CEE3DMP through the

```
ras_dumpoptions_ledumpoptions
```

variable. For more information about CEE3DMP and its options, see z/OS Language Environment Programming Reference, SA22-7562..

**Default:** 3

**Example:**

```
ras_dumpoptions_dumptype=2
```

**ras\_dumpoptions\_ledumpoptions= options**

Specifies dump options to be used with a CEE3DMP. If you want more than one option, separate each option with a blank. Specifies dump options to be used with a CEE3DMP. If you want more than one option, separate each option with a blank.

This WebSphere variable is used only when you specify

```
ras_dumpoptions_dumptype=4
```

. For an explanation of CEE3DMP and valid dump options, see z/OS Language Environment Programming Reference, SA22-7562.

**Rule:** The maximum length of the option string on this environment variable is 255. If the option string is longer than 255, you receive message BBOM0011W and the CEE3DMP dump options are set to

```
THREAD(ALL) BLOCKS
```

.

**Default:**

```
THREAD(ALL) BLOCKS
```

**Example:**

```
ras_dumpoptions_ledumpoptions=NOTRACEBACK NOFILES
```

## Controlling behavior through timeout values

Timeout variables let you control the amount of time you will allow for various requests to complete. Some of these variables map to internal variable names. The internal variable names are provided here to aid you with debugging.

### HTTP Transport timeout variables

#### ConnectionIOTimeout

Sets a maximum amount of time, in seconds, that the J2EE server will wait for the complete HTTP request to arrive. Set this variable for each of the HTTP transport definitions on the server. You will need to set this variable for both SSL transport and non-SSL transport. The J2EE server starts the timer after the connection has been established, and cancels the connection if a complete request does not arrive within the specified maximum time limit. Specifying a value of zero disables the time out function.

**Default:** 10 seconds

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers > server > Web Container > HTTP Transport>host>Custom Properties.**

**Internal variable name (for debugging purposes):** If you are debugging a problem in SSL-enabled transport, look for the internal variable name *protocol\_https\_timeout\_input* in the was.env file or the JES job log. If you are debugging a problem in non-SSL transport, look for the internal variable name *protocol\_http\_timeout\_input* in the was.env file or the JES job log.

**Example:** ConnectionIOTimeout=15

### ConnectionResponseTimeout

Sets a maximum amount of time, in seconds, that the J2EE server will wait for an application component to respond to an HTTP request. Set this variable for each of the HTTP transport definitions on the server. You will need to set this variable for both SSL transport and non-SSL transport. If the response is not received within the specified length of time, the servant (region) fails with ABEND EC3 and RC=04130007. Setting this timer prevents client applications from waiting for a response from an application component that might be deadlocked, looping, or encountering other processing problems related to:

- JSP compiles ( javac/jit compiles)
- XML parser
- jaspr

**Default:** 120 seconds

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers > server > Web Container > HTTP Transport >host> Custom Properties.**

**Internal variable name (for debugging purposes):** If you are debugging a problem in SSL-enabled transport, look for the internal variable name *protocol\_https\_timeout\_output* in the was.env file or the JES job log. If you are debugging a problem in non-SSL transport, look for the internal variable name *protocol\_http\_timeout\_output* in the was.env file or the JES job log.

**Example:** ConnectionResponseTimeout=150

### ConnectionKeepAliveTimeout

Specifies the time, in seconds, that the J2EE server will wait for a subsequent request from an HTTP client on a persistent connection. If another request is not received from the same client within this time limit, the connection is closed.

**Default:** 30 seconds

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers > server > Web Container > HTTP Transport >host> Custom Properties.**

**Internal variable name (for debugging purposes):** If you are debugging a problem in SSL-enabled transport, look for the internal variable name *protocol\_https\_timeout\_persistentSession* in the was.env file or the JES job log. If you are debugging a problem in non-SSL transport, look for the internal variable name *protocol\_http\_timeout\_persistentSession* in the was.env file or the JES job log.

**Example:** ConnectionKeepAliveTimeout=40

## ORB Service Advanced Settings

### ORB listener keep alive

This variable defines the value, in seconds, provided to TCP/IP on the *SOCK\_TCP\_KEEPALIVE* option for the IIOP listener. The function of this

option is to verify if idle sessions are still valid by polling the client TCP/IP stack. If the client does not respond, the session is closed. If the client goes away without notifying the server, it would unnecessarily leave the session active on the server side. Use this option to clean up these unnecessary sessions.

- If the environment variable is not set, the TCP/IP option is not set.
- Setting the *SOCK\_TCP\_KEEPALIVE* option generates network traffic on idle sessions, which can be undesirable.

**Default:** 0

**How to specify:** Specify this setting in the administrative console using the path **Application Servers > server > ORB Service > Advanced Settings**

**Internal variable name (for debugging purposes):** Look for the internal variable name *protocol\_iiop\_server\_session\_keepalive* in the was.env file or the JES job log.

### **ORB SSL listener keep alive**

This variable defines the value, in seconds, provided to TCP/IP on the *SOCK\_TCP\_KEEPALIVE* option for the IIOP listener. The function of this option is to verify if idle sessions are still valid by polling the client TCP/IP stack. If the client does not respond, the session is closed. If the client goes away without notifying the server, it would unnecessarily leave the session active on the server side. Use this option to clean up these unnecessary sessions.

- If the environment variable is not set, the TCP/IP option is not set.
- Setting the *SOCK\_TCP\_KEEPALIVE* option generates network traffic on idle sessions, which can be undesirable.

**Default:** 0

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers > server > ORB Service > Advanced Settings**

**Internal variable name (for debugging purposes):** Look for the internal variable name *protocol\_iiop\_server\_session\_keepalive\_ssl* in the was.env file or the JES job log.

### **WLM timeout**

Specifies the maximum amount of time, in seconds, that WebSphere Application Server for z/OS will wait for IIOP requests to complete. This time limit includes:

- Time during which the IIOP request waits on the WLM queue until being dispatched to a servant (region), and
- Time during which an application component, running in the servant, processes the request and generates a response.

The server generates a failure response if this processing does not complete within the specified time.

**Note:** This variable setting does not apply for HTTP requests; for that type of work, the value specified through the *ConnectionResponseTimeout* variable controls the time allowed for dispatching work to a servant (region).

**Default:** 300 seconds

**How to specify:** Specify this setting in the administrative console using the path **Application Servers > server > ORB Service > Advanced Settings**.



**Internal variable name (for debugging purposes):** Look for the internal variable name *control\_region\_wlm\_dispatch\_timeout* in the was.env file or the JES job log.

**Example:** WLM timeout=600

### Request timeout

Specifies the maximum time, in tenths of seconds, that the client will wait for the response to a client request. This variable is the only timeout available for remote method dispatches made by clients only, not by application components within the servant region. Because the sysplex TCP/IP that runs through the coupling facility does not always tell the client when the other end of the socket is gone, clients can wait indefinitely for a response unless you set this variable. Setting *Request timeout* ensures that the client gets a response within the specified time, even if the response is a COMM\_FAILURE exception.

**Default:** 0 (unlimited), which means no time-out value is set

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers > server > ORB Service**

**Internal variable name (for debugging purposes):** Look for the internal variable name *protocol\_iiop\_local\_timeout* in the was.env file or the JES job log.

**Example:** Request timeout=20 sets the time to 2 seconds

### Transaction Service timeout variables

#### Total Transaction Lifetime Timeout

Specifies the maximum amount of time, in seconds, that the J2EE server will wait for an application transaction to complete. This default amount of time is given to the application transaction if it does not set its own timeout value through the *current->set\_timeout* method. The timer in the controller (region) starts when the container starts a global transaction. If the application transaction is not committed or rolled back within the specified time, the controller abends servant (region) in which the application component is running, with ABEND EC3 reason code 04130002 or 04130005.

The maximum value is 2147483 seconds (24.85 days). You should not use a null or 0 value.

**Default:** 120 seconds

**How to specify:** Specify this setting in the administrative console using the path **Application Servers > server > Transaction Service**

**Internal variable name (for debugging purposes):** Look for the internal variable name *transaction\_defaultTimeout* in the was.env file or the JES job log.

#### Maximum Transaction Timeout

Specifies the maximum amount of time, in seconds, that your installation will allow an application to specify for its transactions to complete. If an application assigns a greater amount of time through the *current->set\_timeout* method, the J2EE server overrides the application setting to the value specified for the *Maximum Transaction Timeout* variable. If the application does not set its own time-out value through the *current->set\_timeout* method, the default value set through the *Maximum Transaction Timeout* variable applies.

The maximum value is 2147483 seconds (24.85 days). You should not use a null or 0 value.

**Default:** 300 seconds

**How to specify:** Specify this setting in the administrative console using the path **Application Servers** > *server* > **Transaction Service**

**Internal variable name (for debugging purposes):** Look for the internal variable name *transaction\_maximumTimeout* in the was.env file or the JES job log.

#### **transaction\_recoveryTimeout**

Specifies the time, in minutes, that this controller (region) uses to attempt to resolve in-doubt transactions before issuing a write-to-operator-with-reply (WTOR) message to the console, requesting whether it should:

- Stop trying to resolve in-doubt transactions,
- Write transaction-related information to the job log or hard copy log, and
- Terminate.

If the operator replies that recovery should continue, the controller (region) will attempt recovery for the specified amount of time before re-issuing the WTOR message. Once all the transactions are resolved, the control region terminates. This variable applies only to controllers in peer restart and recovery mode.

**Default:** 15 minutes

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers** > *server* > **Custom Properties**

**Internal variable name (for debugging purposes):** The internal variable name is the same as the external variable name. Look for *transaction\_recoveryTimeout* in the was.env file or the JES job log.

**Example:** `transaction_recoveryTimeout=7`

### **Server Custom Properties**

#### **control\_region\_mdb\_request\_timeout**

Specifies the time, in seconds, that the server will wait for a message driven bean (MDB) request to receive a response. If the response is not received within the specified amount of time, the server will clean up the MDB request. Set this value to zero to disable the function.

**Default:** 120

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers** > *server* > **Custom Properties**.

**Internal variable name (for debugging purposes):** The internal variable name is the same as the external variable name. Look for *control\_region\_mdb\_request\_timeout* in the was.env file or the JES job log.

**Example:** `control_region_mdb_request_timeout=180`

#### **protocol\_http\_timeout\_output\_recovery**

Controls the recovery action taken on timeouts for requests received over the HTTP transport. Specifying "SERVANT" allows for the termination of server regions when timeouts occur. If an HTTP request is under dispatch in a server region when its timeout value is reached, the server region terminates with an ABEND EC3 RSN=04130007. The HTTP request and socket are then cleaned up. A setting of "SESSION" only cleans up the HTTP request and socket. No attempt is made to disrupt the execution of a dispatched HTTP request within a server region. Be careful using this setting as it may lead to a loss of resources if the dispatched HTTP request loops or hangs.

**Default:** SERVANT

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers > server > Custom Properties**.

**Internal variable name (for debugging purposes):** The internal variable name is the same as the external variable name. Look for *protocol\_http\_timeout\_output\_recovery* in the was.env file or the JES job log.

**Example:** protocol\_http\_timeout\_output\_recovery=SERVANT

#### **protocol\_https\_timeout\_output\_recovery**

Controls the recovery action taken on timeouts for requests received over the HTTP SSL transport. Specifying "SERVANT" allows for the termination of server regions when timeouts occur. If an HTTP SSL request is under dispatch in a server region when its timeout value is reached, the server region terminates with an ABEND EC3 RSN=04130007. The HTTP SSL request and socket are then cleaned up. A setting of "SESSION" only cleans up the HTTP SSL request and socket. No attempt is made to disrupt the execution of a dispatched HTTP SSL request within a server region. Be careful using this setting as it may lead to a loss of resources if the dispatched HTTP SSL request loops or hangs.

**Default:** SERVANT

**How to specify:** Specify this custom property in the administrative console using the path **Application Servers > server > Custom Properties**

**Internal variable name (for debugging purposes):** The internal variable name is the same as the external variable name. Look for *protocol\_https\_timeout\_output\_recovery* in the was.env file or the JES job log.

**Example:** protocol\_https\_timeout\_output\_recovery=SESSION

### **Secure Sockets Layer Configuration Repertoires**

#### **V3 Timeout**

Specifies the length of time that a browser can reuse a System SSL Version 3 session ID without renegotiating encryption keys with the server. The repertoires that you define for a server require the same V3 timeout value.

**Default:** 100

**How to specify:** Specify this custom property in the administrative console using the path **Security > SSL Application Servers > New SSL Repertoire**

**Internal variable name (for debugging purposes):** The following SSL Configuration Repertoire timeout variables are set internally when you define your SSL repertoires:

- com\_ibm\_CSI\_perform\_ssl\_sys\_v3\_timeout
- com\_ibm\_CSI\_claim\_ssl\_sys\_v3\_timeout
- com\_ibm\_HTTP\_claim\_ssl\_sys\_v3\_timeout
- com\_ibm\_DAEMON\_claim\_ssl\_sys\_v3\_timeout

Look for these internal variables in the was.env file or the JES job log.

---

## Debugging WebSphere Application Server applications

In order to debug your application, you must use your application development tool (such as WebSphere Studio Application Developer) to create a Java project or a project with a Java nature. You must then import the program that you want to debug into the project. By following the steps below, you can import the WebSphere Application Server examples into a Java project.

There are two debugging styles available:

- **Step-by-step** debugging mode prompts you whenever the server calls a method on a Web object. A dialog lets you step into the method or skip it. In the dialog, you can turn off step-by-step mode when you are finished using it.
- **Breakpoints** debugging mode lets you debug specific parts of programs. Add breakpoints to the part of the code that you must debug and run the program until one of the breakpoints is encountered.

Breakpoints actually work with both styles of debugging. Step-by-step mode just lets you see which Web objects are being called without having to set up breakpoints ahead of time.

You need not import an entire program into your project. However, if you do not import all of your program into the project, some of the source might not compile. You can still debug the project. Most features of the debugger work, including breakpoints, stepping, and viewing and modifying variables. You must import any source that you want to set breakpoints in.

The inspect and display features in the source view do not work if the source has build errors. These features let you select an expression in the source view and evaluate it.

1. Create a Java Project by opening the New Project dialog.
2. Select **Java** from the left side of the dialog and **Java Project** in the right side of the dialog.
3. Click **Next** and then specify a name for the project (such as WASExamples).
4. Press **Finish** to create the project.
5. Select the new project, choose **File > Import > File System**, then **Next** to open the import file system dialog.
6. Select the directory Browse pushbutton and go to the following directory:  
*installation\_root*\installedApps\*node\_name*\DefaultApplication.ear\DefaultWebApplication.war.
7. Select the checkbox next to DefaultWebApplication.war in the left side of the Import dialog and then click **Finish**. This will import the JavaServer Pages files and Java source for the examples into your project.
8. Add any JAR files needed to build to the Java Build Path. To do this, select **Properties** from the right-click menu. Choose the Java Build Path node and then select the Libraries tab. Use the Add External JARs pushbutton to add the following JAR files:
  - *installation\_root*\installedApps\*node\_name*\DefaultApplication.ear\Increment.jar.  
Once you have added this JAR file, select it and use the **Attach Source** pushbutton to attach Increment.jar as the source - Increment.jar contains both the source and class files.
  - *installation\_root*\lib\j2ee.jar
  - *installation\_root*\lib\pagelist.jar

- `installation_root\lib\webcontainer.jar`

Click **OK** when you have added all of the JARs.

9. You can set some breakpoints in the source at this time if you like, however, it is not necessary as step-by-step mode will prompt you whenever the server calls a method on a Web object. Step-by-step mode is explained in more detail below.
10. To start debugging, you need to start the WebSphere Application Server in debug mode and make note of the JVM debug port. The default value of the JVM debug port is 7777.
11. Once the server is started, switch to the debug perspective by selecting **Window > Open Perspective > Debug**. You can also enable the debug launch in the Java Perspective by choosing **Window > Customize Perspective** and selecting the **Debug** and **Launch** checkboxes in the **Other** category.
12. Select the workbench toolbar **Debug** pushbutton and then select **WebSphere Application Server Debug** from the list of launch configurations. Click the **New** pushbutton to create a new configuration.
13. Give your configuration a name and select the project to debug (your new WASExamples project). Change the port number if you did not start the server on the default port (7777).
14. Click **Debug** to start debugging.
15. Load one of the examples in your browser (for example, `http://localhost:9080/hitcount`).

## Attaching WebSphere Studio Application Developer to a remote debug session

The steps below describe how to attach WebSphere Studio Application Developer to a remote debug session on WebSphere Application Server for z/OS. Remote debugging can prove useful when the program you are debugging behaves differently on a z/OS system than on your local system.

1. Enable the debug engine on WebSphere Application Server for z/OS using the administrative console. See "Debugging Service details" on page 89.
2. Import the the java source code you wish to debug into WebSphere Studio Application Developer and set breakpoints. See the *WebSphere Studio Application Developer information center* at <http://publib.boulder.ibm.com/infocenter/wsphelp/> for instructions on setting breakpoints.
3. Open a WebSphere Studio Application Developer debug perspective and create a debug session configuration.
4. Attach WebSphere Studio Application Developer to the WebSphere Application Server for z/OS debug runtime. See "Connecting to a remote VM with the remote Java application launch configuration" in the *WebSphere Studio Application Developer information center* at <http://publib.boulder.ibm.com/infocenter/wsphelp/>
5. Execute the java code in WebSphere Application Server for z/OS to hit the breakpoints set in WebSphere Studio Application Developer.
6. Use WebSphere Studio Application Developer debugger controls and features to debug the application.

## Unit testing with DB2

The steps below describe how to setup a unit test environment that would allow you to develop and unit test code with DB2 z/OS to support Container Managed Persistence (CMP) development and access DB2 test data that resides on z/OS.

1. Configure DB2 Distributed Data Facility on z/OS to allow remote TCP/IP connections from your WebSphere Studio Application Developer workstation. See the DB2 Information Center for information on DDF.
2. Install the DB2 Client Configuration Assistant on the workstation where WebSphere Studio Application Developer is installed. The DB2 Client Configuration Assistant is shipped with DB2.
3. Use the DB2 Client Configuration Assistant to define a DB2 alias.
4. Use the DB2 alias you defined to access the DB2 subsystem on z/OS using the DB2 Distributed Data Facility (DDF).





---

## Chapter 5. Viewing diagnostic information

The following topics provide information about specific sources of diagnostic data, and the tools or resources you might need to view or work with that data.

| Type of diagnostic tools or data: | Notes and instructions for use appear in:                            |
|-----------------------------------|----------------------------------------------------------------------|
| CEEDUMPs                          | "Viewing CEEDUMPs in the job log"                                    |
| SVC dumps                         | Viewing SVC dumps                                                    |
| CTRACE and JRas data              | Viewing CTRACE and JRas data through IPCS                            |
| Error log data                    | Viewing error log contents through the Log Browse Utility (BBORBLOG) |
| z/OS display command              | Using the z/OS display command                                       |
| Java minor codes                  | Converting Java minor codes                                          |
| SYSPRINT                          | Viewing SYSPRINT output, STDOUT and STDERR streams in an HFS File    |
| Message routing                   | Message routing                                                      |

---

### Viewing CEEDUMPs in the job log

An error caught by LE or the Java run-time can result in the production of a CEEDUMP, which is written to a separate CEEDUMP specification in the job log. To view the dump contents, select the CEEDUMP portion of the output for the address space. The 'Traceback' section at the beginning of the dump can be very helpful.

---

### Viewing SVC dumps

A SVC dump is a core dump initiated by the operating system generally when a programming exception occurs. SVC dump processing stores data in dump data sets that you pre-allocate, or that the system allocates automatically as needed.

Alternatively, you can initiate an SVC dump through the MVS console, to gather diagnostic data for a 'hang' condition, for example. SVC dumps that you initiate this way are called console dumps.

One example of an abend that could occur is the EC3 abend. WebSphere Application Server for z/OS requests an SVC dump when a controller terminates a servant (region) with an EC3 abend when timeout conditions occur.

- Your installation can set parmlib options that determine what to dump, eliminate duplicate dumps, and so on. WebSphere Application Server for z/OS provides a dump parmlib sample in  
SBBOJCL(BBODMCCB)

- The standard SDATA expected in a SVC dump:  
SDATA=(ALLNUC,CSA,GRSQ,LPA,LSQA,PSA,RGN,SQA,SUM,SWA,TRT),end

- If you cannot find an SVC dump for a specific abend, your installation might be using Dump analysis and elimination (DAE) to suppress the dump. If this is the case, you can change DAE to let the dump be taken or set a SLIP on the specific abend for a particular job name if the timeout is consistently happening. For further information, see:
  - z/OS MVS Diagnosis: Tools and Service Aids, GA22-7589 for details about using DAE.
  - z/OS MVS System Commands, SA22-7627 for details about the SLIP command, which controls SLIP (serviceability level indication processing), a diagnostic aid that intercepts or traps certain system events and specifies what action to take. Using the SLIP command, you can set, modify, and delete SLIP traps.
- When you initiate a console dump:
  - When you want an SVC dump of a servant region, also request a dump of the servant's controller region.
  - Unless you suspect a particular servant region as the source of a problem, dump the controller region and all of its servant regions.
- If syslog contains a message indicating that the maxspace limit was reached for this dump, the SVC dump might be a partial one that might not contain the data you need to diagnose the timeout. This limit means that the data set used for SVC dump is not large enough, and you have to change the size to capture a complete dump.
- To view CEEDUMP contents within the SVC dump, use the IPCS verbexit LEDATA, with the CEEDUMP or NTHREADS options, to format and analyze Language Environment control blocks. For additional information, see z/OS Language Environment Debugging Guide, GA22-756 for instructions for using IPCS to format and analyze CEEDUMP contents.

See z/OS MVS Diagnosis: Tools and Service Aids, GA22-7589 for additional information about SVC dumps.

---

## Viewing CTRACE and JRas data through IPCS

Once activated, the WebSphere Application Server for z/OS always writes trace data into memory buffers. The number and size of these buffers is controlled using WebSphere variables. You can get this trace data from a dump, which may be taken by the system or requested by the operator through DUMP or SLIP commands.

To view messages or application trace data from Component Trace, you must use the interactive problem control system (IPCS) to format the data. The source of the trace data can be a dump data set or a trace data set, and the command to use would be

```
IPCS CTRACE
```

. You can also use the

```
IPCS CTRACE
```

command to merge multiple trace entities together such as multiple WebSphere Application Server for z/OS address space traces, OMVS, and TCPIP.

## Steps for using the IPCS dialog to format CTRACE data

When setting up IPCS, your installation may customize IPCS for its users. IBM recommends providing access to the IPCS dialog through an ISPF panel. If your

installation has not customized IPCS as recommended, you need to start the IPCS dialog. See *z/OS MVS IPCS User's Guide, SA22-7596* to find out how to start the IPCS dialog.

Perform the following steps to use the IPCS dialog to format application trace data:

1. From the IPCS Primary Option Menu panel, select option 6 ( *COMMAND* ).
2. On the IPCS Subcommand Entry panel:
  - a. Issue the *SETDEF* subcommand to determine the default values for routing displays.
  - b. Enter the *CTRACE* command, with the following required parameters:

```
CTRACE COMP(cell_short_name
)
```

where *cell\_short\_name* is the value specified through the ISPF customization dialog to identify the location of server configuration files. The name must be 8 or fewer characters and all uppercase.

**Note:** If you were interested in only JRAS data, you would enter the following:

```
CTRACE COMP(cell_short_name
)USEREXIT(JRAS)
```

Specify additional parameters as necessary.

**Example:** To direct trace data to the terminal only, you would append the

```
NOPRINT
```

and

```
TERMINAL
```

parameters to the

```
CTRACE
```

command.

**Tip:** For a complete list of CTRACE command parameters, see *z/OS MVS IPCS Commands, SA22-7594*.

3. View your application's data, basing the method you choose on which one is appropriate for the location of the data:

| If you directed output to the... | Then use the...             |
|----------------------------------|-----------------------------|
| IPCS print data set (IPCSPRNT)   | ISPF/PDF Browse option      |
| Terminal                         | Dump Display Reporter panel |

**Tips:** To navigate through the trace data on the Dump Display Reporter panel, use the commands and PF keys listed in *z/OS MVS IPCS User's Guide, SA22-7596*.

CTRACE enables you to view multiple traces together with the trace data from the various sources intermixed based on the time stamp. See *z/OS MVS IPCS Commands, SA22-7594*, for specifics on using this MERGE subcommand.

## Finding the subname for IPCS CTRACE

If the trace data set is an SVC dump, the trace subname must also be specified. This subname is the aggregation of the address space's jobname with its ASID (address space identifier), in printable hexadecimal. An easy way to determine the subname is to query CTRACE for the data using the following IPCS subcommand:

```
CTRACE QUERY DSN('dump.data.set')
```

Once you get the subname you can view the WebSphere Application Server for z/OS trace data with the following IPCS subcommand:

```
CTRACE COMP(cell_short_name) SUB((subname)) FULL DSN('dump.data.set')
```

where *cell\_short\_name* is the value specified through the ISPF customization dialog to identify the location of server configuration files. The name must be 8 or fewer characters and all uppercase.

**Note:** The *subname* parameter is optional for only the trace data set. It is required when viewing the trace data using the dump data set.

## Steps for using IPCS in batch mode to format CTRACE data

You must create an IPCS dump directory before you can use IPCS in batch mode. When setting up IPCS, your installation may customize IPCS for its users. This customization can include modifying the IBM-supplied BLSCDDIR CLIST with default values for creating an IPCS dump directory.

To view messages or application trace data from Component Trace, you must use the interactive problem control system (IPCS) to format the data. Using IPCS in batch mode is the easiest method of formatting data, especially if you do not have much experience with using IPCS, TSO/E and ISPF. Through batch mode, you can use IPCS to format trace data and write it to an MVS data set. Optionally, you may copy the contents of that data set into an HFS file for viewing.

When your installation has modified the BLSCDDIR CLIST the steps outlined herein will create an IPCS dump directory.

1. Decide on a fully-qualified data set name for the directory.
2. From the TSO/E command prompt, enter the

```
BLSCDDIR
```

command, specifying the data set name.

For example, to create a dump directory named IBMUSER.DDIR, enter:

```
%blscddir dsn('ibmuser.ddir')
```

If your installation has not customized IPCS, you might need to alter other BLSCDDIR CLIST parameters. See the z/OS MVS IPCS User's Guide, SA22-7596 and z/OS MVS IPCS Commands, SA22-7594 for more details about using the BLSCDDIR CLIST to create a dump directory.

Perform the following steps to use IPCS in batch mode to format application trace data:

1. Create a file and copy the following sample JCL into it. This JCL invokes IPCS to extract and format JRAS trace data and write it into an MVS data set, and then uses the

```
TSO/E OPUT
```

command to copy the formatted data from the MVS data set into an HFS file.

```
//IBMUSERX JOB ,
// CLASS=J,NOTIFY=&SYSUID,MSGCLASS=H
//IPCS EXEC PGM=IKJEFT01,REGION=4096K,DYNAMNBR=50
//IPCSDDIR DD DSN=IBMUSER.DDIR,DISP=SHR
//IPCSDOC DD SYSOUT=H
//JRASTRC DD DSN=IBMUSER.CB390.CTRACE,DISP=SHR
//IPCSPRNT DD DSN=IBMUSER.IPCS.OUT,DISP=OLD
```

```

//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
IPCS
DROPDUMP DDNAME(JRASTRC)
PROFILE LINESIZE(80)PAGESIZE(99999999)
SETDEF NOCONFIRM
CTRACE COMP(SYSBBOSS) DDNAME(JRASTRC) FULL PRINT +
 NOTERMIAL
DROPDUMP DDNAME(JRASTRC)
END
/*
//OPUT EXEC PGM=IKJEFT01,REGION=4096K,DYNAMNBR=50
//SYSTSPRT DD SYSOUT=*
//SYSTSIN DD *
oput 'ibmuser.ipcs.out' '/u/ibmuser/ipcs/jrastrace.txt' TEXT
/*

```

2. Edit the sample JCL to replace *IBMUSER.DDIR* with the data set name that you used for the IPCS dump directory you created.

- a. Use the *PAGESIZE* parameter on the *PROFILE* statement only if you do not want to print the output data set.
- b. You may replace the HFS file name with the name of an existing HFS file, but you do not have to do so. The

*OPUT*

command processing will create a new HFS file, if the one specified does not exist, and grants read and write access to that file for your user ID only.

If you do specify an existing HFS file, the

*OPUT*

command processing will write over any data that is already in that file. If you want to know more about the

*OPUT*

command, see the z/OS UNIX System Services Command Reference, SA22-7802.

- c. Change the data set name specified on the *JRASTRC DD* in the example to the name of the data set containing the *CTRACE* data.
- d. Change the name of the MVS data set on both the *JRASTRC DD* statement and the *OPUT* command in the *SYSTSIN* stream, as necessary. The formatted output of the *JRAS CTRACE* data is first written to the MVS data set specified by the

*IPCSPRNT DD*

statement and then (optionally) copied to the HFS data set. You must either pre-allocate this data set, or change the sample JCL to allocate the data set.

This data set should have a record format of VBA and a record length of 133.

3. Submit the JCL to start the IPCS batch job.

Once you are done you can use a UNIX editor, such as *vi*, to view your trace data in the HFS file. If you want to know more about the UNIX editors, see z/OS UNIX System Services User's Guide, SA22-7801.

*CTRACE* enables you to view multiple traces together with the trace data from the various sources intermixed based on the time stamp. See z/OS MVS IPCS Commands, SA22-7594, for specifics on using this

*MERGE*

subcommand.

### Sample JCL to display WebSphere for z/OS trace data

The following sample shows JCL that displays WebSphere for z/OS trace data.

**Note:** The JCL uses an IPCS dump directory (in VSAM data set userid.DUMP.DIR) that must be allocated before you run the JCL. See z/OS MVS IPCS Commands, SA22-7594 , for information about initializing a dump directory.

```
//SHOWTRC JOB <job card info>
//JOBLIB DD DISP=SHR,DSN=BBO.SBBOMIG
// DD DISP=SHR,DSN=SYS1.MIGLIB
//PRINTIT EXEC PGM=IKJEFT01,REGION=OM
//IPCSDDIR DD DISP=(OLD,KEEP),DSN=userid.DUMP.DIR
//IPCSPARM DD DISP=SHR,DSN=SYS1.PARMLIB
//SYSTSPRT DD SYSOUT=*
//IPCSTOC DD SYSOUT=*
//IPCSPRNT DD SYSOUT=*
//*-----
//SYSTSIN DD *
IPCS NOPARM
TRACE COMP(SYSBBOSS) SUB((subname)) FULL DSN('dump.data.set')
/*
```

The following example shows JCL that displays WebSphere for z/OS trace data for multiple address spaces.

```
//SHOWTRC2 JOB <job card info>

//JOBLIB DD DISP=SHR,DSN=BBO.SBBOMIG
// DD DISP=SHR,DSN=SYS1.MIGLIB
//PRINTIT EXEC PGM=IKJEFT01,REGION=OM
//IPCSDDIR DD DISP=(OLD,KEEP),DSN=userid.DUMP.DIR
//IPCSPARM DD DISP=SHR,DSN=SYS1.PARMLIB
//SYSTSPRT DD SYSOUT=*
//IPCSTOC DD SYSOUT=*
//IPCSPRNT DD SYSOUT=*
//*-----
//SYSTSIN DD *
IPCS NOPARM
MERGE
TRACE COMP(SYSBBOSS) SUB((subname)) FULL DSN('dump.data.set')
TRACE COMP(SYSBBOSS) SUB((subname2)) FULL DSN('dump.data.set')
MERGEEND
/*
```

---

## Viewing error log contents through the Log Browse Utility (BBORBLOG)

You can use the Log Browse Utility (BBORBLOG) to view the error log stream. If you need to look at the WebSphere Application Server for z/OS error log stream, use ISPF option 6 to enter the command:

1. Use ISPF option 6 to enter the proper command.  
ex 'BBO.SBBOEXEC(BBORBLOG)' BBO.BOSSXXXX  
where the log-stream name is BBO.BOSSXXXX
2. The space allocation and the unit for the allocation are contained within the REXX code. If you keep a large amount of trace data, the allocation must be made larger.
3. The WebSphere Application Server for z/OS provides an ISPF REXX EXEC named BBORBLOG, that allows you to browse the error log stream.

#### 4. Save the output.

When you use the BBORBLOG browser, it creates a data set with your user ID followed by the log stream name. You should rename it if you wish to save your browser output. The contents of the current view of the log stream will remain until the stream reaches its retention date. The next time you invoke the browser, however, the current view of the log stream will be deleted (because it uses the same data set name). The previous data will exist in another record (not the current view) until its retention date.

Use the following information to determine viewing of the error log:

- By default, the macro formats the error records to fit a 3270 display.
- Timestamps are in Greenwich Mean Time (GMT) unless changed by setting the WebSphere Application Server variable `ras_time_local` to 1.
- Message BBOJ0051I, which appears in the job output, can help correlate error-log entries to the proper job output.

## Using the log browse utility (BBORBLOG)

The browser takes two parameters:

| Parameter       | Description                                                                                                                                                                                                                                                            |
|-----------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| log stream name | The name of the log stream. See the job messages for the name of the log stream.                                                                                                                                                                                       |
| format option   | <b>80</b> The default. The log stream record will be formatted on a <code>lrecl</code> length of 80 characters. Additional lines will be wrapped.<br><b>NOFORMAT</b> Turns off formatting. The error log message appears as one log message string in the browse file. |

View the error log stream output using the BBORBLOG browser. To invoke the browser, go to ISPF option 6 and enter:

```
'BBO.SBBOEXEC(BBORBLOG)' BBO.BOSSXXXX format option
```

**Note:** In this example, BBORBLOG resides in BBO.SBBOEXEC.

The browser creates a browse data set named "userid.stream\_name", which contains the contents of the log stream. When the browser is executed, it:

1. Allocates a data set called `userid.stream_name`, which overwrites any duplicate data sets.
2. Populates the data set with the contents of the log stream.
3. Puts the user in "browse" mode on the data set.

**Important:** Each time BBORBLOG is invoked a static file is created which overwrites the existing file. In order to refresh the file, it is necessary to re-issue BBORBLOG

There are three valid ways (three separate commands to use) to invoke the browser. We will illustrate each of these using the following example:

**Example:** If the BBORBLOG member was in a data set named BBO.SBBOEXEC, then you would issue one of the following depending on your chosen format option:

```
ex 'BBO.SBBOEXEC(BBORBLOG)' BBO.BOSSXXXX
ex 'BBO.SBBOEXEC(BBORBLOG)' BBO.BOSSXXXX 80
ex 'BBO.SBBOEXEC(BBORBLOG)' BBO.BOSSXXXX NOFORMAT
```



**Tip:** It is easier to invoke the browser if you add the target library (in our example, BBO.SBBOEXEC) to the SYSEXEC concatenation of the user logon procedure during the WebSphere Application Server for z/OS installation. If you concatenate the library during installation, you do not have to specify the library containing the browser REXX exec- you only need to specify BBORBLOG.

## Error log stream record output

There are two error log stream records that we will look at:

- Server logstream
- CERR of a server.

**Note:** The numbers to the left of each sample were added to specify lines-they will not be in the actual output.

### Sample output from a server logstream:

```

1 | 2000/06/01 16:01:06.683 01 SYSTEM=SY1 SERVER=BBOASR1A JobName=BBOASR1S
2 | ASID=0X0033 PID=0X0100003C TID=0X24F858A0 0X000004 c=2.1010030
3 | ./bbooreq.cpp+4437 ... BBOU0013W The function
4 | make_user_exception(IIOP_protocolArea*)+4437 raised a user exception
5 | CosNaming::NamingContext::NotFound.
```

The log stream record output fields from stream BBO.BOSSXXXX are:

*Table 3. Parts table for a server logstream record output*

| Component                                                 | Description                                                    |
|-----------------------------------------------------------|----------------------------------------------------------------|
| line 1: 2000/06/01 16:01:06.683 01                        | Date / timestamp / 2-digit record version number               |
| line 1: SYSTEM=SY1                                        | System name                                                    |
| line 1: SERVER=BBOASR1A                                   | Server name                                                    |
| line 1: JobName=BBOASR1S                                  | Jobname                                                        |
| line 2: ASID=0X0033                                       | ASID (address space identifier)                                |
| line 2: PID=0X0100003C                                    | PID (Process ID)                                               |
| line 2: TID=0X24F858A0 0X000004                           | TID (Thread ID)                                                |
| line 2: c=2.1010030                                       | Request correlation information                                |
| line 3: ./bbooreq.cpp+4437                                | File name & line                                               |
| line 3: BBOU0013W                                         | Log message number                                             |
| line 3: The function. ..                                  | Log message                                                    |
| lines 4-5: make_user_exception...<br>CosNaming::Naming... | Continuation lines: Continuation of the Log Stream log message |

**Note:** Each field is delimited by a blank.

### Sample output from CERR of a server:

```

1 | BossLog: { 0017} 2000/06/01 15:58:25.557 01 SYSTEM=SY1 SERVER=BBOASR1A
2 | PID=0X0100003C TID=0X24F82920 00000000 c=3.C5D02
3 | ./bboiroot.cpp+1195 ... BBOU0012W The function IRootHomeImpl::findHome(
4 | const char*)+1195 received CORBA system exception CORBA::INTERNAL.
5 | Error code is C9C21200.
```

The CERR job message output fields are:

Table 4. Parts table for a CERR record output

| Component                                                                                              | Description                                                    |
|--------------------------------------------------------------------------------------------------------|----------------------------------------------------------------|
| line 1: BossLog: { 0017}                                                                               | BossLog: {entry number}                                        |
| line 1: 2000/06/01 15:58:25.557 01                                                                     | Date / timestamp / 2-digit record version number               |
| line 1: SYSTEM=SY1                                                                                     | System name                                                    |
| line 1: SERVER=BBOASR1A                                                                                | Server name                                                    |
| line 2: PID=0X0100003C                                                                                 | PID (Process ID)                                               |
| line 2: TID=0X24F82920 00000000                                                                        | TID (Thread ID)                                                |
| line 2: c=3.C5D02                                                                                      | Request correlation information                                |
| line 3: . /bboiroot.cpp+1195                                                                           | File name & line                                               |
| line 3: BBOU0012W                                                                                      | Log message number                                             |
| line 3: The function IRootHomeImpl::find. ..                                                           | Log message                                                    |
| lines 4-5: const char*))+1195 received CORBA system exception CORBA::INTERNAL. Error code is C9C21200. | Continuation lines: Continuation lines of the CERR job message |

- Each field is delimited by a blank.
- The CERR format is found in SYSOUT, not the logger.

### Saving your BBORBLOG browser output

When you use the BBORBLOG browser, it creates a data set with your user ID followed by the log stream name. You should rename it if you wish to save your browser output. The contents of the current view of the log stream will remain until the stream reaches its retention date. The next time you invoke the browser, however, the current view of the log stream will be deleted (because it uses the same data set name). The previous data will exist in another record (not the current view) until its retention date.

---

## Using the z/OS display command

You may use either the WebSphere Administrative console or the z/OS MVS console to accomplish many operations tasks related to WebSphere Application Server for z/OS servers. Entering the z/OS display or modify commands through the MVS console can provide information or perform tasks that are useful for diagnosing problems.

In particular, the z/OS display command provides parameters that allow you to display information about the following:

- Servers
- Servant regions
- Sessions
- Trace settings
- JVM heap statistics
- work
- error log

---

## Converting Java minor codes

Occasionally, Java will take an WebSphere Application Server for z/OS error code (C9C2xxxx in hexadecimal) and convert it to a very large negative number. If you get a very large negative number, try converting it back to hexadecimal to find the correct code.

To convert the error codes back to hexadecimal:

- Add  $2^{32}$  to the negative number and convert it into hexadecimal. This can be done using the OMVS command

```
bc
```

**Example:** Suppose you get the error code "910022649":

1. Under OMVS, type the command:

```
bc
```

2. then type:

```
obase=16
2^32 - 910022649
quit
```

- The bc program displays C9C22807, which is the hex value that you should look up.

---

## Viewing SYSPRINT output, STDOUT and STDERR streams in an HFS File

If you are familiar with UNIX or NT environments, you may be reluctant to use the facilities of SDSF (or IOF) to view the SYSPRINT output from your Application Server regions. In the past, you may have used a familiar editor (such as vi) in a TELNET session to view the STDOUT and STDERR streams. These STDOUT and STDERR streams have been redirected to SYSPRINT, and it is possible to redirect them to files in the HFS for viewing using an HFS editor.

The JCL below for an Application Server region uses this facility. A new SET statement has been added to point to the LOGPATH directory, and the SYSPRINT DD statement has been changed to point to a file in the LOGPATH/servername directory, in this case named: was.log.d &LYMMDD.t&LHHMSS. The extra period (.) between the date and time variables is not a typographical error, but rather an idiosyncrasy of JCL syntax that is necessary to terminate the first variable. &LYMMDD will be replaced with the local date in YYMMDD format and &LHHMSS will be replaced by the local time in HHMMSS format. Using the local date and time ensures a unique file for each instance of the Application Server region that is started. The PATHMODE subparameter sets the file mode to 775 and the PATHOPTS subparameter OWRONLY opens the file for WRITE access and the sub-parameter OCREAT indicates that if the file does not already exist, create it.

```
//TSTST1S PROC IWMSSNM='TSTST1S',PARMS='-ORBsrvname '
// SET CBCONFIG='/WebSphere390/TSDCONF'
// SET LOGPATH='/WebSphere390/logs' <<----- Added to set path
// SET RELPATH='controlinfo/envfile'
//WSDAS1S EXEC PGM=BBOSR,REGION=0M,TIME=NOLIMIT,
// PARM='/ &PARMS &IWMSSNM'
//BBOENV DD PATH='&CBCONFIG/&RELPATH/&SYSPLEX/&IWMSSNM/current.env'
//CEEDUMP DD SYSOUT=*
//SYSOUT DD SYSOUT=*
//*SYSPRINT DD SYSOUT=* Added the following 3 lines for SYSPRINT
```

```
//SYSPRINT DD PATHMODE=(SIRWXU,SIRWXG,SIROTH), <<< 775
// PATHOPTS=(OWRONLY,OCREAT), <<<
// PATH='&LOGPATH/TSTST1S/was.log.d&LYMMDD..t&LHHMSS'
```

In the previous example, if the file specified by the SYSPRINT DD statement is created on May 2, 2003 at 1:30:35 PM, the PATH parameter will resolve to: /WebSphere390/logs/TSTST1S/was.log.d030503.t133035 which will be a unique file for this execution of this server instance.

---

## Message routing

The following is a summary of how messages get routed in WebSphere Application Server Version 5:

- WTO messages issued by WebSphere Application Server during initialization are sent to hardcopy, but most can be routed to the data set specified by ras\_default\_msg\_dd (see Setting Output destinations and characteristics).
- The Java "Audit" messages are also sent to hardcopy, but can be routed to the data set ras\_hardcopy\_msg\_dd. (see Setting Output destinations and characteristics).
- Trace error, service, and fatal messages are sent to the error log specified by the ras\_log\_logstreamName. Otherwise, they go to CERR (SYSOUT). Some may also go to hardcopy. You must set the ras\_log\_logstreamName environment variable to the error log stream name. To set ras\_log\_logstreamName, on the administrative console select Environment -> Manage WebSphere Variables -> New.
- Early error messages will go to STDERR (SYSOUT) until WebSphere Application Server connects to the log stream. A WTO (BBOO0153I) is issued telling you how many messages went to STDERR before you connected to the log stream.
- Trace messages (such as BUFFER, SYSPRINT, and TRCFILE DD) are also routed to ras\_trace\_outputLocation.
- System.out.println, Tr.Sysout, stdout (cout), and client output go to SYSPRINT.(see Viewing SYSPRINT output, STDOUT and STDERR streams in an HFS File
- System.out.printerr, Tr.Syserr, and stderr (cerr) go to SYSOUT.

To use these new message routing variables, you must do two things:

1. Add these parameters to the server definitions using the Administrative Console under Environment -> Manage WebSphere Variables:
  - **ras\_default\_msg\_dd =DEFALTDD**
  - **ras\_hardcopy\_msg\_dd =HRDCPYDD**

You can set these variables for individual control and servant processes, but it is easier to set them in the Environment variables for the entire cell. For the Daemon, you must prefix them with "DAEMON\_" and set them at the cell level:

- **DAEMON\_ras\_default\_msg\_dd =DEFALTDD**
- **DAEMON\_ras\_hardcopy\_msg\_dd =HRDCPYDD**

Update the included &Z members in PROCLIB to add these new DD statements

```
/* Output DDs
//CEEDUMP DD SYSOUT=*,SPIN=UNALLOC,FREE=CLOSE
//SYSOUT DD SYSOUT=*,SPIN=UNALLOC,FREE=CLOSE
```

```
//SYSPRINT DD SYSOUT=*,SPIN=UNALLOC,FREE=CLOSE
//DEFALTDD DD SYSOUT=*,SPIN=UNALLOC,FREE=CLOSE
//HRDCPYDD DD SYSOUT=*,SPIN=UNALLOC,FREE=CLOSE
```

**Note:**

- If you specify the new environment variables, but do not specify the DD cards in the procedure, you will not get an error message indicating that the DD cards are missing, and the tracing output will not be written anywhere.
- If you try to direct the output for multiple streams to the same DD, such as setting both **ras\_default\_msg\_dd** and **ras\_hardcopy\_msg\_dd** to DEFALTDD (or to SYSPRINT,) then the allocation will fail and output will be sent to the default location (JOBLOG/SYSLOG).

For example, these DD files are used to segregate the messages and keep almost all of them off the hardcopy console

1. (SYSLOG.) JESMSGGLG - a few start-up and shut-down messages
2. JESYSMSG - MVS allocation and deallocation messages
3. SYSOUT - a few start-up and shut-down messages
4. SYSPRINT - a few start-up and shut-down messages
5. HRDCPYDD - audit messages that would normally go to SYSLOG
6. DEFALTDD - informational messages that would normally go to SYSLOG

## Using the Error Dump and Cleanup interface

The Error Dump and Cleanup (BBORLEXT) interface exists to call WebSphere Application Server for z/OS in a recovery environment to allow it to request a dump and to clean up its resources.

The interface will:

- Save the function and DLL names of the failing z/OS component into the SDWA.
- Determine whether or not to issue an SDUMP, if relevant to the time-of-failure environment.
- Clean up z/OS internal structures and connections.

**Program requirements:** This interface **must** be called from within a WebSphere Application Server for z/OS location service daemon, controller (region), or servant (region). There are no restrictions against in which recovery environment, such as an ESTAE or FRR routine, the caller must reside.

### General information

|                     |                                        |
|---------------------|----------------------------------------|
| Interface:          | BALR to BBORLEXT                       |
| Address of routine: | (ECVT+'234'x)+'20'x                    |
| Address mode:       | AMODE 31, RMODE any                    |
| State:              | Allow problem program state, task mode |
| Cross memory mode:  | PASN=HASN=SASN (non-cross memory)      |
| Return codes:       | No return codes                        |

|           |                                                                                                            |
|-----------|------------------------------------------------------------------------------------------------------------|
| Function: | Clean-up various WebSphere for z/OS resources and possibly issue an SVC dump for the current address space |
|-----------|------------------------------------------------------------------------------------------------------------|

### Input register information

The contents of the registers are as follows:

|    |                                              |
|----|----------------------------------------------|
| 1  | Contains the address of the SDWA             |
| 14 | Contains the return address                  |
| 15 | Contains the entry point address of BBORLEXT |

### Output register information

When control returns to the caller, the contents of the registers are as follows:

|      |                                       |
|------|---------------------------------------|
| 0-1  | Used as a work register by the system |
| 2-14 | Unchanged                             |
| 15   | Used as a work register by the system |

**Note:** Some callers depend on register contents remaining the same before and after issuing a service. If the system changes the contents of registers on which the caller depends, the caller must save them before issuing the service and restore them after the system returns control.

**Note:** A dump will not occur for X22 abends or for certain reason codes from 0D6, 052, 067, CC3, and DC3 abends. There may also be other error conditions that will not create a dump.

#### Example:

Example Here is an example of how to call this routine in assembler:

```
LA 1,SDWA Load SDWA@ in Reg 1
L 15,(0,16) Load CVT address
L 15,140(,15) Load ECVT address
L 15,564(,15) Load address of z/OS structure
L 15,32(,15) Load address of z/OS routine
BALR 14,15 Invoke z/OS routine
```

---

## Viewing the output data set

The data set should be viewed using a program that can display record type 120.

The Java SMF Record Interpreter is provided in the form of a jar file named `bbomsmfv.jar`. To use it from the z/OS or OS/390 UNIX environment:

1. Verify that the `JAVA_HOME` environment variable refers to the current java installation. `JAVA_HOME=../usr/bin/java/J1.3` This should be at least Java 1.3 since this release is the first to implicitly contain the necessary record support needed by the interpreter.

2. Download the SMF Record Interpreter from the WebSphere for z/OS web site at: [http://www.ibm.com/software/webservers/appserv/zos\\_os390/index.html](http://www.ibm.com/software/webservers/appserv/zos_os390/index.html)
3. Copy the file `bbomsmfv.jar` to your tools directory. Be sure that any edits made to the file in the future are made to both copies of the file, or just execute from the installation directory in the first place.
4. To interpret SMF data from a cataloged z/OS or OS/390 sequential file named "USER.SMFDATA" (which was previously created using the IFASMFDP utility as described above), execute: `java -cp bbomsmfv.jar com.ibm.ws390.sm.smfview.Interpreter "USER.SMFDATA"` It is implicit in the java command parameterization that your current working directory is the tools directory. If this is not the case, you will receive a `NoClassDefFoundError` on `com.ibm.ws390.sm.smfview.Interpreter`. Java doesn't generate a diagnostic when it doesn't find `bbomsmfv.jar` in the current directory.

The SMF ViewTool has been successfully installed and invoked when you do not receive any Java error messages after the invocation and the Browser output is shown on the screen.

## Example of SMF Browser output

Example of SMF Browser output

The SMF Browser available on the WebSphere for z/OS download site is able to display record type 120. To download the SMF Browser go to: <http://www6.software.ibm.com/dl/websphere20/zosos390-p>. For further information on the SMF Browser, download the browser package and read the associated documentation.

The following example shows sample output from the SMF Browser. The example features subtype 7 and subtype 8, in that order.

```
Record#: 14;
 Type: 120; Size: 820; Date: Fri Nov 23 04:54:17 EST 2001;
 SystemID: SY1; SubsystemID: WAS; Flag: 94;
 Subtype: 7 (WEB CONTAINER ACTIVITY);
Triplets: 4;
Triplet #: 1; offset: 76; length: 32; count: 1;
Triplet #: 2; offset: 108; length: 140; count: 1;
Triplet #: 3; offset: 264; length: 556; count: 1;
Triplet #: 1; Type: ProductSection;
 Version: 1; Codeset: Unicode; Endian: 1; TimeStampFormat: 1 (S390STCK64);
 IndexOfThisRecord: 1; Total # records: 1; Total # triplets: 4;
Triplet #: 2; Type: WebContainerActivitySection;
 HostName : PLEX1;
 ServerName : BBOASR4;
 ServerInstanceName: BBOASR4A;
 WlmEnclaveToken * 00000020 00000242 -----
 * ^... * p1047
 ActivityID * b6c7a7b7 14e9bc85 000000b0 00000007
 * 0926306b -----
 *,..... *Cp1047
 ActivityStartTime * b6c7a7b7 14e9bc85 40404040 40404040 *
 ActivityStopTime * b6c7a7b7 53a8a645 40404040 40404040 *
Triplet #: 3; Type: HttpSessionManagerActivitySection;
 # http sessions created: 0; # http sessions invalidated: 0;
 # http sessions active: 0;
Average session life time: 0 [sec*10**-3];
Triplet #: 4; Type: WebApplicationActivitySection;
 Name: PolicyIVP-localhost_1;
 # Servlets: 1;
 Triplet #: 4.1; offset: 272; length: 284; count: 1;
 Triplet #: 4.1; Type: ServletActivitySection;
```



```

Name: SimpleFileServlet;
ResponseTime: 48 [sec*10**-3];
errors: 0;
Loaded by this request: 0;
Loaded since (raw): ea54948e0d;
Loaded since: Thu Nov 22 10:02:49 EST 2001;
Record#: 72;
Type: 120; Size: 1744; Date: Fri Nov 23 05:01:02 EST 2001;
SystemID: SY1; SubsystemID: WAS; Flag: 94;
Subtype: 8 (WEB CONTAINER INTERVAL);
Triplets: 4;
Triplet #: 1; offset: 76; length: 32; count: 1;
Triplet #: 2; offset: 108; length: 112; count: 1;
Triplet #: 3; offset: 264; length: 1480; count: 1;
Triplet #: 1; Type: ProductSection;
 Version: 1; Codeset: Unicode; Endian: 1; TimeStampFormat: 1 (S390STCK64);
 IndexOfThisRecord: 1; Total # records: 1; Total # triplets: 4;
Triplet #: 2; Type: WebContainerIntervalSection;
 HostName : PLEX1;
 ServerName : BBOASR4;
 ServerInstanceName: BBOASR4A;
 SampleStartTime * b6c7a6fd 655c0604 40404040 40404040 *
 SampleStopTime * b6c7a939 9a0e614c 40404040 40404040 *
Triplet #: 3; Type: HttpSessionManagerIntervalSection;
 http sessions #created: 1; #invalidated: 0;
 http sessions #active: 0; Min #active: 0; Max #active: 0;
 Average session life time: 0;
 Average session invalidate time: 0;
 http sessions #finalized: 0; #tracked: 0;
 http sessions #min live: 0; #max live: 0;
Triplet #: 4; Type: WebApplicationIntervalSection;
 Name: PolicyIVP-localhost_1;
 # Servlets loaded: 0;
 # Servlets: 4;
 Triplet #: 4.1; offset: 312; length: 292; count: 1;
 Triplet #: 4.2; offset: 604; length: 292; count: 1;
 Triplet #: 4.3; offset: 896; length: 292; count: 1;
 Triplet #: 4.4; offset: 1188; length: 292; count: 1;
 Triplet #: 4.1; Type: ServletIntervalSection;
 Name: SimpleFileServlet;
 # requests: 6;
 AverageResponseTime: 764 [sec*10**-3];
 MinimumResponseTime: 18 [sec*10**-3];
 MaximumResponseTime: 4133 [sec*10**-3];
 # errors: 0;
 Loaded since (raw): ea54948e0d;
 Loaded since: Thu Nov 22 10:02:49 EST 2001;
 Triplet #: 4.2; Type: ServletIntervalSection;
 Name: Was40Ivp;
 # requests: 4;
 AverageResponseTime: 4664 [sec*10**-3];
 MinimumResponseTime: 1584 [sec*10**-3];
 MaximumResponseTime: 12572 [sec*10**-3];
 # errors: 0;
 Loaded since (raw): ea58a1509e;
 Loaded since: Fri Nov 23 04:55:14 EST 2001;
 Triplet #: 4.3; Type: ServletIntervalSection;
 Name: /cebit.jsp;
 # requests: 1;
 AverageResponseTime: 204 [sec*10**-3];
 MinimumResponseTime: 204 [sec*10**-3];
 MaximumResponseTime: 204 [sec*10**-3];
 # errors: 0;
 Loaded since (raw): ea58a24a69;
 Loaded since: Fri Nov 23 04:56:18 EST 2001;
 Triplet #: 4.4; Type: ServletIntervalSection;
 Name: JSP 1.1 Processor;

```

```
requests: 1;
AverageResponseTime: 482 [sec*10**-3];
MinimumResponseTime: 482 [sec*10**-3];
MaximumResponseTime: 482 [sec*10**-3];
errors: 0;
Loaded since (raw): ea54948b66;
Loaded since: Thu Nov 22 10:02:48 EST 2001;
```

---

## Chapter 6. Automation and recovery scenarios and guidelines

The following section provides information on how to monitor and recover WebSphere Application Server for z/OS and the subsystems it uses. It provides startup, shutdown, and recovery procedures and scenarios. It also tells you how to determine if the subsystems are up or down, and tells you where to find more information.

---

### APPC automation and recovery scenarios

Table 5. APPC automation and recovery scenarios

| Task                                                             | APPC automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup                                                          | APPC should be started before WebSphere Application Server for z/OS. In theory, WebSphere Application Server for z/OS <i>could</i> be started before APPC, but only as long as no objects get dispatched in containers that have an IMS APPC LRMI associated with them. If APPC is not up before WebSphere Application Server for z/OS, and you want to use an APPC connector to talk to IMS, then you will have no connectivity. APPC/MVS does not have to be up for CICS. APPC does not have to be started after VTAM.                                                                                                                              |
| Shutdown                                                         | Reverse the startup procedure. Shutdown WebSphere Application Server for z/OS, APPC, then VTAM.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| Handling in-flight or indoubt transactions if there is a failure | If you are using APPC for communications and it fails, do the following: <ol style="list-style-type: none"><li>1. Shutdown all servers with APPC connectivity.</li><li>2. Restart APPC (if it totally failed).</li><li>3. Restart the WebSphere Application Server for z/OS server.</li></ol> <p><b>Note:</b> APPC will resynchronize itself. If your transaction is indoubt, IMS waits until you restart APPC. IMS relies on RRS for recovery. RRS will resolve transactions that are in doubt by handshaking with every subsystem it was communicating with before it went down. If you are using CICS, note that CICS has its own coordinator.</p> |
| How to determine if APPC is running                              | Issue the DISPLAY APPC,LU,ALL command. If APPC is not active, it will say so. In addition, the status of the logical units used by WebSphere Application Server for z/OS and/or IMS should be active or no APPC work will be successful.                                                                                                                                                                                                                                                                                                                                                                                                              |
| What happens to WebSphere for z/OS if APPC goes down?            | Any objects attempting to use the IMS APPC PAA will not work. The server region running on behalf of the container attempting to use APPC will likely get a C9C24C05 error, indicating that an APPC ALLOCATE request was attempted and failed. Additional APPC error diagnostic information that helps to pinpoint the APPC problem is contained in the logs associated with this region.                                                                                                                                                                                                                                                             |
| What happens to other subsystems if APPC goes down?              | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |

Table 5. APPC automation and recovery scenarios (continued)

| Task                           | APPC automation and recovery scenarios                                                                                                                                                 |
|--------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Where to find more information | <ul style="list-style-type: none"> <li>• z/OS MVS Planning: Operations</li> <li>• z/OS MVS Planning: APPC/MVS Management</li> <li>• z/OS MVS Programming: Resource Recovery</li> </ul> |

## WLM automation and recovery scenarios

Table 6. Workload Manager (WLM) automation and recovery scenarios

| Task                                                                                             | WLM automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                    |
|--------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup                                                                                          | WLM is automatically started by z/OS or OS/390 when you IPL your system. You do not have to start it.                                                                                                                                                                                                                                                                                                                                                    |
| Shutdown                                                                                         | You cannot shutdown WLM.                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| Handling in-flight and indoubt transactions if there is a failure                                | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| How to determine if WLM is running                                                               | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| What happens to WebSphere for z/OS if WLM goes down?                                             | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| What happens to other subsystems if WLM goes down?                                               | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| How to handle a catastrophic failure of the WebSphere Application Server for z/OS server regions | <p>Following a catastrophic failure of the WebSphere Application Server for z/OS server regions, you can use one of the following resume commands, depending on your application environment type:</p> <ul style="list-style-type: none"> <li>• Static application environment:<br/> <code>v wlm,applenv=applenvname, resume</code> </li> <li>• Dynamic application environment:<br/> <code>v wlm,applenv=applenvname, resume,options</code> </li> </ul> |
| Where to find more information                                                                   | <ul style="list-style-type: none"> <li>• z/OS MVS Planning: Workload Management</li> <li>• z/OS MVS Programming: Workload Management Services</li> </ul>                                                                                                                                                                                                                                                                                                 |

## RACF automation and recovery scenarios

Table 7. RACF automation and recovery scenarios

| Task                                                              | RACF automation and recovery scenarios                       |
|-------------------------------------------------------------------|--------------------------------------------------------------|
| Startup                                                           | If it is installed, RACF is started as a part of IPL.        |
| Shutdown                                                          | RACF is not shutdown.                                        |
| Handling in-flight and indoubt transactions if there is a failure | Not applicable                                               |
| How to determine if RACF is running                               | Use the RACF SETROPTS command to display the status of RACF. |

Table 7. RACF automation and recovery scenarios (continued)

| Task                                                  | RACF automation and recovery scenarios                                                                                                                                                    |
|-------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| What happens to WebSphere for z/OS if RACF goes down? | RACF goes into fail safe mode. This means that for every resource that is accessed, the operator is asked to verify if it is okay. In general, the system is IPLed if this occurs.        |
| What happens to other subsystems if RACF goes down?   | It depends on the subsystem and how RACF fails.                                                                                                                                           |
| Where to find more information                        | <ul style="list-style-type: none"> <li>• <i>z/OS Security Server RACF System Programmer's Guide</i></li> <li>• <i>z/OS Security Server RACF Security Administrator's Guide</i></li> </ul> |

## RRS automation and recovery scenarios

Table 8. RRS automation and recovery scenarios

| Task     | RRS automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|----------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup  | <p>Ensure System Logger has been started before RRS.<br/> <b>Note:</b> RRS will display error messages indicating that System Logger must be started first if you try to start RRS without starting System Logger. Ensure RRS is started before WebSphere for z/OS. RRS does not start by itself. RRS will start automatically only if it was registered with the Automatic Restart Manager (ARM) and if ARM is running. To start RRS, issue the start command:</p> <pre>start atrrrs,sub=master</pre> <p><b>Note:</b> RRS doesn't restart itself if you issue the cancel command, so you need to restart it manually if it was canceled or if ARM isn't running.</p> |
| Shutdown | <p>Shutdown RRS in the reverse order that you started RRS. Shutdown WebSphere Application Server for z/OS, then RRS, followed by System Logger. There is no controlled way to bring down RRS. The best approach is:</p> <ol style="list-style-type: none"> <li>1. Quiesce WebSphere Application Server for z/OS.</li> <li>2. Shutdown WebSphere Application Server for z/OS.</li> <li>3. Cancel RRS.</li> </ol> <p><b>Note:</b> You may want to bring down the DB2 you are using for WebSphere Application Server for z/OS before canceling RRS.</p> <p>To cancel RRS, issue the command:</p> <pre>setrrs cancel</pre>                                                |

Table 8. RRS automation and recovery scenarios (continued)

| Task                                                              | RRS automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Handling in-flight and indoubt transactions if there is a failure | <p>Refer to the RRS system management panels to display in-flight and resolve indoubt transactions. You can display the resource managers on the RM panels in RRS, display all units of recovery (UR), filter the URs, and then resolve the indoubts. You cannot resolve in-flights. You can display all RRS-managed transactions.</p> <p>If you are using the IMS Connector for Java, this process applies only if IMS Connector for Java, IMS Connect, and the IMS subsystem have been configured locally on the same z/OS or OS/390 system image on which the WebSphere for z/OS J2EE server runs. The local configuration is the only configuration in which IMS Connector for Java runs as an RRS-transactional connector.</p> |
| How to determine if RRS is running                                | <p>Use the display command:</p> <pre>d a,atrrs</pre> <p>atrrs is the name of the default RRS procedure shipped with WebSphere Application Server for z/OS. Use the procedure name that you use to start RRS. The address space comes from the procedure.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| What happens to WebSphere for z/OS if RRS goes down?              | <p>RRS is a required subsystem, so WebSphere Application Server for z/OS will not run without it. If RRS goes down, WebSphere Application Server for z/OS will get fatal errors. You need to get RRS started, then restart WebSphere Application Server for z/OS.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| What happens to other subsystems if RRS goes down?                | <p>RRS is the z/OS or OS/390 transaction monitor. If you cancel RRS, you will have problems with any subsystems using it (for example, WebSphere Application Server for z/OS, DB2, IMS). You should understand the implications before you cancel RRS.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
| Where to find more information                                    | <p><i>z/OS MVS Programming: Resource Recovery</i></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |

## UNIX System Services automation and recovery scenarios

Table 9. UNIX System Services automation and recovery scenarios

| Task                                                                  | UNIX System Services automation and recovery scenarios                                                                                  |
|-----------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Startup                                                               | <p>UNIX System Services is a permanent component of MVS and is started automatically at IPL time.</p>                                   |
| Shutdown                                                              | <p>UNIX System Services does not support a shutdown capability, so it is always available.</p>                                          |
| Handling in-flight or indoubt transactions if there is a failure      | <p>The only data that could be considered transactional in nature is data stored in the HFS.</p>                                        |
| How to determine if UNIX System Services is running                   | <p>UNIX System Services is always available as long as the system is up and running.</p>                                                |
| What happens to WebSphere for z/OS if UNIX System Services goes down? | <p>If UNIX System Services fails, the system must be re-IPLed. WebSphere Application Server for z/OS will get errors and terminate.</p> |

Table 9. UNIX System Services automation and recovery scenarios (continued)

| Task                                                                | UNIX System Services automation and recovery scenarios      |
|---------------------------------------------------------------------|-------------------------------------------------------------|
| What happens to other subsystems if UNIX System Services goes down? | If UNIX System Services fails, the system must be re-IPLed. |
| Where to find more information                                      | <i>z/OS UNIX System Services Planning</i>                   |

## TCP/IP automation and recovery scenarios

Table 10. TCP/IP automation and recovery scenarios

| Task                                                                       | TCP/IP automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                             |
|----------------------------------------------------------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup                                                                    | TCP/IP must be up before starting WebSphere Application Server for z/OS.                                                                                                                                                                                                                                                                                                             |
| Shutdown                                                                   | Shutdown WebSphere Application Server for z/OS before shutting down TCP/IP.                                                                                                                                                                                                                                                                                                          |
| Handling inflight or indoubt transactions if there is a failure            | Methods in flight will have their transactions rolled back when the attempt to send a response to the method fails. Other transactions will wait for a timeout.                                                                                                                                                                                                                      |
| How to determine if TCP/IP is running                                      | Use the display command looking for the TCP/IP procedure.                                                                                                                                                                                                                                                                                                                            |
| What happens to WebSphere Application Server for z/OS if TCP/IP goes down? | If TCP/IP goes down, then WebSphere Application Server for z/OS on the system must be restarted. You will get an SVC dump because the socket layer was destroyed.                                                                                                                                                                                                                    |
| What happens to other subsystems if TCP/IP goes down?                      | If TCP/IP goes down, sessions break and transactions react as described above.<br><b>Note:</b> WebSphere Application Server for z/OS will not be able to recognize that TCP/IP is back up. Therefore, WebSphere Application Server for z/OS must be restarted.<br><b>Note:</b> If TCP/IP goes down, you should recycle LDAP before restarting WebSphere Application Server for z/OS. |

## DB2 automation and recovery scenarios

Table 11. DB2 automation and recovery scenarios

| Task     | DB2 automation and recovery scenarios                                                     |
|----------|-------------------------------------------------------------------------------------------|
| Startup  | DB2 is started after RRS but before LDAP, NFS, and WebSphere Application Server for z/OS. |
| Shutdown | Reverse of startup sequence.                                                              |



Table 11. DB2 automation and recovery scenarios (continued)

| Task                                                                    | DB2 automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|-------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Handling in-flight or indoubt transactions if there is a failure        | <p>Use the RRS panels to resolve. See <i>z/OS MVS Programming: Resource Recovery</i>. The RRS panels are the preferred way to resolve DB2 indoubts because they allow you to view all resource managers that have an interest in the transaction. However, you can also use DB2 to resolve indoubts. You can issue the command:</p> <pre>DISPLAY THREAD(*) TYPE(INDOUBT)</pre> <p>to display DB2 information about the indoubt threads it knows about (if there are too many, you can go into S.LOG to view the information). This display will give you a DB2 identifier called a "nid". Copy the nid and paste it into this command:</p> <pre>-RECOVER INDOUBT (RRSAF) ACTION(COMMIT) NID(B1D379D17ED6CF900000009401010000)</pre> <p>where the nid is the one that you cut from the display command. You can issue this command to roll back the transaction:</p> <pre>-RECOVER INDOUBT (RRSAF) ACTION(ABORT) NID(B1D379D17ED6CF900000009401010000)</pre> |
| How to determine if DB2 is running                                      | Use the display command to display the DB2 address space.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| What happens to WebSphere Application Server for z/OS if DB2 goes down? | WebSphere Application Server for z/OS needs DB2 to run, so it abends if DB2 goes down. You need to restart the LDAP server, then restart DB2, and then restart WebSphere Application Server for z/OS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| What happens to other subsystems if DB2 goes down?                      | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| Where to find more information                                          | See the DB2 books at the following Internet location: <a href="http://www.ibm.com/servers/eserver/zseries/zos/">http://www.ibm.com/servers/eserver/zseries/zos/</a>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |

## CICS automation and recovery scenarios

Table 12. CICS automation and recovery scenarios

| Task     | CICS automation and recovery scenarios                                                                                                                                                                                                                  |
|----------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup  | CICS and any required CICS products, such as CICS Transaction Gateway, need to be properly installed, initialized, and started before any work flows to a CICS-enabled WebSphere Application Server for z/OS application control server region are run. |
| Shutdown | Shutdown the WebSphere Application Server for z/OS application control region that uses CICS as a backing store, then shut down the CICS service.                                                                                                       |

Table 12. CICS automation and recovery scenarios (continued)

| Task                                                                     | CICS automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|--------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Handling in-flight or indoubt transactions if there is a failure         | If there is an error during processing, both CICS and WebSphere Application Server for z/OS rely on the underlying RRS subsystem to handle all rollback notifications to the registered interests. In the case of inflight transactions, RRS will notify all participants that a rollback is required, and normal rollback processing will occur in each registered party. In the case of indoubt transactions, it may be necessary to recycle the WebSphere Application Server for z/OS Application Control/Server region to release any pending transaction in CICS. |
| How to determine if CICS is running                                      | This is installation dependent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| What happens to CICS if WebSphere Application Server for z/OS goes down? | Should WebSphere Application Server for z/OS happen to go down, one of two situations could occur: <ol style="list-style-type: none"> <li>1. If WebSphere Application Server for z/OS and CICS are currently engaged in a unit of work, then RRS processing as described above would occur and it may be necessary to recycle the application control server regions to release pending transactional work in CICS.</li> <li>2. If WebSphere Application Server for z/OS and CICS are not currently engaged in a unit of work, CICS is not affected.</li> </ol>        |
| What happens to other subsystems if CICS goes down?                      | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Where to find more information                                           | <i>CICS Operations and Utilities Guide</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |

## IMS automation and recovery scenarios

Table 13. IMS automation and recovery scenarios

| Task                                                             | IMS automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup                                                          | IMS and any required IMS products, such as IMS Connect, need to be properly installed, initialized, and started before any work flows to an IMS-enabled WebSphere for z/OS application control server region are run.                                                                                                                                                                                                                                                                                                                                                        |
| Shutdown                                                         | Shutdown the WebSphere Application Server for z/OS application Control Region which uses IMS as a backing store, then shutdown the IMS service                                                                                                                                                                                                                                                                                                                                                                                                                               |
| Handling in-flight or indoubt transactions if there is a failure | If there is an error during processing, both IMS and WebSphere Application Server for z/OS rely on the underlying RRS subsystem to handle all rollback notifications to the registered interests. In the case of inflight transactions, RRS will notify all participants that a rollback is required and normal rollback processing will occur in each registered party. In the case of indoubt transactions, it may be necessary to recycle the WebSphere Application Server for z/OS Application Control/Server region to release any pending transaction in the IMS MPRs. |

Table 13. IMS automation and recovery scenarios (continued)

| Task                                                                    | IMS automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| How to determine if IMS is running                                      | This is installation-dependent.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| What happens to IMS if WebSphere Application Server for z/OS goes down? | Should WebSphere Application Server for z/OS happen to go down, one of two situations could occur: <ol style="list-style-type: none"> <li>1. If WebSphere Application Server for z/OS and IMS are currently engaged in a unit of work, then RRS processing as described above would occur and it may be necessary to recycle the application control server regions to release pending transactional work in the IMS MPR.</li> <li>2. If WebSphere Application Server for z/OS and IMS are not currently engaged in a unit of work, IMS is not affected.</li> </ol> |
| What happens to other subsystems if IMS goes down?                      | Not applicable                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| Where to find more information                                          | <i>IMS/ESA Operator's Reference</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## LDAP automation and recovery scenarios

Table 14. LDAP automation and recovery scenarios

| Task                                                             | LDAP automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                   |
|------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup                                                          | LDAP, as used by WebSphere Application Server for z/OS, is completely run within the WebSphere Application Server for z/OS address spaces using something called "the local backend." This support takes the front side of the LDAP client APIs and the backend database implementation and runs them completely inside the WebSphere Application Server for z/OS Naming Server and Interface Repository. For Naming and IR, OMVS and DB2 must be up <i>before</i> Naming and IR. To run the LDAP server, TCPIP, OMVS, and DB2 must all be up before the LDAP server. <p><b>Note:</b> There are two LDAP modes supported:</p> <ol style="list-style-type: none"> <li>1. Local LDAP backend.</li> <li>2. Remote LDAP Server. The WebSphere Application Server for z/OS environment has to be set up correspondingly, and DB2, TCP/IP, and the remote server have to be up and running before WebSphere Application Server for z/OS is started.</li> </ol> |
| Shutdown                                                         | Shutdown Naming and IR, then OMVS and DB2. For the LDAP server, shutdown the LDAP server, then TCPIP and DB2, and then OMVS.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| Handling in-flight or indoubt transactions if there is a failure | If there is a failure during processing, Naming and IR rely on RRS to issue a rollback directly to DB2 and, as a result, any work done by the LDAP code is rolled back along with it. For the LDAP server, AUTOCOMMIT is set to NO, causing any error to ROLLBACK for that transaction. This ensures the atomicity characteristic of LDAP operations.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

Table 14. LDAP automation and recovery scenarios (continued)

| Task                                                                     | LDAP automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|--------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| How to determine if LDAP is running                                      | In the case of WebSphere Application Server for z/OS, if Naming and IR are operating, then LDAP is operating. In the case of the LDAP server, and a started task is used for the LDAP server, use the SDSF to see if the started task is running. Examine the output log for the started task to see if any error messages were displayed. Alternatively, the LDAPSRCH command (from TSO), or LDAPSEARCH command (from Unix System Services shell) can be used to perform a simple search to verify that the LDAP server is running.                                                                                                                                                                                                                                                                                                                                                           |
| What happens to WebSphere Application Server for z/OS if LDAP goes down? | <ul style="list-style-type: none"> <li>In J2EE server regions, the LDAP server <b>must</b> be active since it is a separate server that no longer runs inside the WebSphere Application Server for z/OS region. Recycle the JNDI, then restart the J2EE application servers.</li> </ul>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| What happens to other subsystems if LDAP goes down?                      | Most z/OS or OS/390 subsystems do not depend on LDAP, but this may change in the future. In the case of accessing LDAP through the LDAP server, there is a way to configure the LDAP server to operate in a sysplex environment such that (using sysplex-enabled DNS) LDAP requests will be sent to the LDAP server in the sysplex that is operating (assuming that there is one). As an alternative, subsystems that want to use LDAP could configure a backup LDAP server to be contacted in case the primary server is not accessible. In this case, the application would assume that it could retrieve all of the same data that it could get from the backup on the primary which would be handled by some replication mechanism. The LDAP server currently supports a master/slave replication mechanism, but you could also try duplicating the sysplex server using DB2 data sharing. |
| Where to find more information                                           | <i>z/OS Security Server LDAP Server Administration and Use</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

## WebSphere Application Server for z/OS (Daemon) automation and recovery scenarios

Table 15. WebSphere Application Server for z/OS automation and recovery scenarios

| Task                                                             | WebSphere for z/OS (daemon) automation and recovery scenarios                                                                                                                                                                                                                                                                                                               |
|------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup                                                          | See the instructions for "Starting the WebSphere for z/OS environment".                                                                                                                                                                                                                                                                                                     |
| Shutdown                                                         | See the instructions for "Stopping the WebSphere for z/OS environment".                                                                                                                                                                                                                                                                                                     |
| Handling in-flight or indoubt transactions if there is a failure | The daemon is a location agent. If the daemon fails during the course of a transaction, locate requests to the daemon will fail. These request failures will be surfaced by the client ORB. If the client is a WebSphere Application Server for z/OS client running in a sysplex, the locate request will be routed to another available daemon in the sysplex, if present. |

Table 15. WebSphere Application Server for z/OS automation and recovery scenarios (continued)

| Task                                                                           | WebSphere for z/OS (daemon) automation and recovery scenarios                                                                                                                                                             |
|--------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| How to determine if the daemon is running                                      | Use the MVS display command.                                                                                                                                                                                              |
| What happens to WebSphere Application Server for z/OS if the daemon goes down? | If the daemon goes down, all WebSphere Application Server for z/OS servers started on the same system will also be terminated.                                                                                            |
| What happens to other subsystems if the daemon goes down?                      | Other subsystems will continue to work fine. As a general rule, if the servers on multiple systems are defined in cluster and there is another one in the sysplex, if the daemon goes down, clients will not be affected. |

## Web server (servlet) automation and recovery scenarios

Table 16. Web server (servlet) automation and recovery scenarios

| Task                                                                               | WebServer automation and recovery scenarios                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Startup                                                                            | Web servers have a relationship with WebSphere Application Server for z/OS only in the sense that a client application program that is written to use WebSphere Application Server for z/OS facilities may be written as a servlet. Any implications for ordering of startup will be introduced by the applications. You probably want to have the WebSphere Application Server for z/OS object servers up and ready before starting the client application that the web server is hosting. |
| Shutdown                                                                           | There are no dependencies from the product code. Similar to most applications, you may want to quiesce the clients prior to taking down the target WebSphere Application Server for z/OS servers. Shut down the web server to stop the port of entry.                                                                                                                                                                                                                                       |
| Handling in-flight or indoubt transactions if there is a failure                   | Since a web server is stateless, there are no in-flight or indoubt transactions.                                                                                                                                                                                                                                                                                                                                                                                                            |
| How to determine if a web server is running                                        | Use the z/OS display commands and viewer tools (SDDF) to monitor the Application Server.                                                                                                                                                                                                                                                                                                                                                                                                    |
| What happens to WebSphere Application Server for z/OS if the web server goes down? | WebSphere Application Server for z/OS requires an IBM HTTP J2EE server web server in order to provide full function servlets. So, if the web server goes down, applications that require a port of entry (like servlets and SOAP) cannot run.                                                                                                                                                                                                                                               |
| What happens to other subsystems if Web Server goes down?                          | HTTP only provides information to other subsystems, so they are unaffected if the web server goes down.                                                                                                                                                                                                                                                                                                                                                                                     |
| Where to find more information                                                     | <i>z/OS HTTP Server Planning, Installing, and Using</i>                                                                                                                                                                                                                                                                                                                                                                                                                                     |

---

## Chapter 7. Preparing for a call to IBM service

When you report a problem to IBM service, you will need to provide as much information as possible to help service personnel quickly resolve the problem. The information you might need to send depends, in part, on the type of problem you have encountered, and includes the following items:

- Job logs for affected address spaces; for example, the controller and any servant regions that the controller terminated
- Job output for affected address spaces, particularly WebSphere Application Server for z/OS messages that are written to the JESMSGLG data set
- The system log (SYSLOG), another source of WebSphere Application Server for z/OS messages
- WebSphere Application Server for z/OS error log
- The system logrec data set or log stream
- CTRACE external writer data sets
- SVC dumps, CEEDUMPs, or other types of dumps produced because of the problem.
- The affected server's environment file, WAS.env, which is located in the HFS:  
`AppServer/config/cells/cellname/nodes/nodename/servers/servername/was.env`

Additionally, IBM service might request you to:

- Provide a description of the circumstances or scenario under which the error occurs.
- Use the  
`VERBEXIT CBDATA`  
subcommand.
- Reset WebSphere variables that are for use only when directed by IBM service.
- Set WebSphere variable values for the location service daemon address space (same as those for servers, with the prefix "DAEMON\_").

---

### Using the IPCS VERBEXIT subcommand to display diagnostic data

The interactive problem control system (IPCS) is a tool that provides formatting and analysis support for dumps and traces produced by WebSphere Application Server for z/OS and the applications that it hosts. IBM service personnel might request that you use the IPCS subcommand VERBEXIT with the CBDATA verb name to display dump information for WebSphere Application Server for z/OS. The CBDATA formatters reside in the SBB0MIG data set, which must be in the link list or LPA.

Entering VERBEXIT CBDATA results in a display of dump contents that can include the following WebSphere Application Server for z/OS structures:

- Global control blocks
- Address space control blocks
- Task control blocks (TCBs)
- ORB control block

Optional parameters control which of these structures are included in the dump display. If you enter VERBEXIT CBDATA without any optional parameters, the dump display includes only global control block contents.

To enter VERBEXIT CBDATA, you may use any of the methods for entering IPCS subcommands on z/OS, as described in z/OS MVS IPCS User's Guide, SA22-7596. Use the following syntax: VERBEXIT CBDATA [ 'parameter [,parameter]...' ]

Valid parameters are:

- **GLOBAL(bgvt-address)**

Formats and displays cell-level global vector data for the specified address space. This display includes the following formatted control blocks:

- BGVT address - z/OS Global Vector table
- ASR Table and ASR Table entries - Active Server Resposity information

- **ASID(asid-number)**

Formats and displays address space information for the daemon, the controller (region), or the servant (region). This display includes the following formatted control blocks:

- BACB - z/OS address space control block
- BTRC,TBUFSET,TBUF - z/OS Component trace control blocks
- BOAM,BOAMX - z/OS BOA control blocks
- ACRW queue - Application Control Region Work element control blocks
- BTCB queues - z/OS control information

Along with ASID(as id-number), IBM service personnel might direct you to specify one of the following parameters, to include additional information in the dump display:

- **BTCB(btcb\_address)**

Formats and displays the specified BTCB and ORB information for the WebSphere Application Server for z/OS TCB.

- **COMMDATA**

Formats and displays session information.

- **CONFIG**

Formats and displays configuration information for the address space.

- **OBJADDR(object\_address)** and **OBJTYPE(object\_type\_ID)**

Formats and displays information for the specified object of the specified type. IBM service personnel will provide the values for you to supply for these parameters.

- **ORBDATA**

Formats and displays ORB information.

- **TCB(tcb\_address)**

Formats and displays request summary information for the specified task.

- **TRACEBACK**

Formats and displays ORB information.

- **SUMMARY**

Summarizes information from some of the other CBDATA optional parameters. For example, for the GLOBAL parameter, specifying SUMMARY produces a list of active servers.

**Example:** Output from the command ip VERBEXIT CBDATA 'ASID(xx ) TCB(yyyyyyy )':

```
command ==> ip VERBX CBDATA 'ASID(xx) TCB(xxxxxxxx)'
***** TOP OF DATA *****
COMPON=WEBSPPHERE Z/OS,COMPID=5655A9801,ISSUER=BBORMCDP,ERRNO=04006006
```

---

```
BBOR0012I Formatting C1ssname
C1ssname: 2BE6947E
+0000 D9859496 A385E685 82C39695 A3818995 |RemoteWebContain|.....|
+0010 859900 |er. |...
```



```

BBOR0012I Formatting MethodNm
 MethodNm: 2BE69472
+0000 88A3A397 998598A4 85A2A300 00000000 |httprequest.....|.....?.....|
BBOR0012I Formatting ComRtInf
 ComRtInf: 2BE69212
+0000 89974081 8484997E F94BF5F6 4BF4F24B |ip addr=9.56.42.|..@....~.K..K..K|
+0010 F1F6F840 979699A3 7EF1F0F8 F500 |168 port=1085. |...@....~.....|
BBOR0026I GMT Time Request was received into CTL region
 TODCLOCK: 00000000
 04/08/2003 12:58:02.926136
BBOR0026I GMT Time Request was Queued to WLM in CTL region
 TODCLOCK: 00000000
 04/08/2003 12:58:02.926263
BBOR0026I GMT Time Request will be Expired (approximated)
 TODCLOCK: 00000000
 04/08/2003 13:08:01.663032
BBOR0026I GMT Time Request was received into SR region
 TODCLOCK: 00000000
 04/08/2003 12:58:02.927729

```

---

## Setting trace controls for IBM service

To view or set your trace settings, use the WebSphere administrative console:

- Click **Environment > Manage WebSphere Variables**.
- On the Configuration Tab check for any of these variables in the name field and observe the variable setting in the value field.
- To change or set a variable, specify the variable in the name field and specify the setting in the value field. You can also describe the setting in the description field on this tab.

### Note:

- If you use any level of tracing, including `ras_trace_defaultTracingLevel=1`, ensure that you set `ras_trace_outputLocation` to `BUFFER`.  
`ras_trace_defaultTracingLevel=1` will write exceptions to the trace log as well as to the ERROR log.

- Set the `ras_trace_BufferCount=4` and `ras_trace_BufferSize=128`.  
This will get 512KB of storage for the trace buffers (the minimum allowed) and reduce memory requirements.

- It is best to trace to `CTRACE`.

If you are tracing to `sysprint` with `ras_trace_defaultTracingLevel=3`, you may experience an almost 100% throughput degradation. If you are tracing to `CTRACE`, however, you may only experience a 15% degradation in throughput.

- Make sure you disable JRAS tracing.

To do this, look for the following lines in the `trace.dat` file pointed to by the JVM properties file:

```
com.ibm.ejs.*=all=disable
```

```
com.ibm.ws390.orb.*=all=disable
```

Ensure that both lines are set to **=disable** or delete the two lines altogether.

**Note:** If `ras_trace_outputLocation` is set, you may be tracing and not know it.

**ras\_trace\_defaultTracingLevel= *n***

Specifies the default tracing level for WebSphere Application Server for z/OS.

Valid values and their meanings are:

|   |                                                         |
|---|---------------------------------------------------------|
| 0 | No tracing                                              |
| 1 | Exception tracing                                       |
| 2 | Basic and exception tracing                             |
| 3 | Detailed tracing, including basic and exception tracing |

Use this variable together with the

`ras_trace_basic`

and

`ras_trace_detail`

variables to set tracing levels for WebSphere Application Server for z/OS subcomponents. Specifies the default tracing level for WebSphere Application Server for z/OS.

**Default:** 1

**Example:**

```
ras_trace_defaultTracingLevel=2
```

**ras\_trace\_basic=*n* | (*n*,...)**

Specifies tracing overrides for particular WebSphere Application Server for z/OS subcomponents.

Subcomponents, specified by numbers, receive basic and exception traces. If IBM service directs you to specify more than one subcomponent, use parentheses and separate the numbers with commas. IBM service provides the subcomponent numbers and their meanings.

Other parts of WebSphere Application Server for z/OS receive tracing as specified on the

```
ras_trace_defaultTracingLevel
```

variable.

**Note:** Valid values for **ras\_trace\_basic** are:

- 0: RAS
- 1: Common Utilities
- 2: COS/Naming
- 3: COMM
- 4: ORB
- 5: IM
- 6: OTS
- 7: Shasta
- 8: Systems Management
- 9: OS/390 Wrappers
- A: Daemon
- B: IR
- C: Test
- D: COS/Query
- E: Security

- F: Externalization
- G: Adapter
- H: Lifecycle
- I: Identity
- J: JRAS (internal tracing--via direction from IBM support)
- K: Reference collections
- L: J2EE

**Default:** (no default value)

**Example:**

```
ras_trace_basic=3
```

**ras\_trace\_detail=n | (n,...)**

Specifies tracing overrides for particular WebSphere Application Server for z/OS subcomponents.

Subcomponents, specified by numbers, receive detailed traces. If IBM service directs you to specify more than one subcomponent, use parentheses and separate the numbers with commas. IBM service provides the subcomponent numbers and their meanings.

Other parts of WebSphere Application Server for z/OS receive tracing as specified on the

```
ras_trace_defaultTracingLevel
```

variable.

**Note:** Valid values for **ras\_trace\_detail** are:

- 0: RAS
- 1: Common Utilities
- 2: COS/Naming
- 3: COMM
- 4: ORB
- 5: IM
- 6: OTS
- 7: Shasta
- 8: Systems Management
- 9: OS/390 Wrappers
- A: Daemon
- B: IR
- C: Test
- D: COS/Query
- E: Security
- F: Externalization
- G: Adapter
- H: Lifecycle
- I: Identity
- J: JRAS (internal tracing--via direction from IBM support)
- K: Reference collections
- L: J2EE

**Default:** (no default value)

**Examples:**

```
ras_trace_detail=3
```

```
ras_trace_detail=(3,4)
```

**ras\_trace\_specific=n | (n,...)**

Specifies tracing overrides for specific WebSphere Application Server for z/OS trace points.

Trace points are specified by 8-digit, hexadecimal numbers. If IBM service directs you to specify more than one trace point, use parentheses and separate the numbers with commas. You also can specify a WebSphere variable name by enclosing the name in single quotes. The value of the WebSphere variable will be handled as if you had specified that value on

`ras_trace_specific`

.

**Default:** (no default value)

**Examples:**

`ras_trace_specific=03004020`

`ras_trace_specific=(03004020,04005010)`

`ras_trace_specific='xyz'`

[where xyz is an environment variable name]

`ras_trace_specific=('xyz','abc',03004021)`

[where xyz and abc are environment variable names]

**ras\_trace\_exclude\_specific=n | (n,...)**

Specifies WebSphere Application Server for z/OS trace points to exclude from tracing activity.

Trace points are specified by 8-digit, hexadecimal numbers. If IBM service directs you to specify more than one trace point, use parentheses and separate the numbers with commas. You also can specify a WebSphere variable name by enclosing the name in single quotes. The value of the WebSphere variable will be handled as if you had specified that value on

`ras_trace_exclude_specific`

.

**Default:** (no default value)

**Examples:**

`ras_trace_exclude_specific=03004020`

`ras_trace_exclude_specific=(03004020,04005010)`

`ras_trace_exclude_specific='xyz'`

[where xyz is an environment variable name]

`ras_trace_exclude_specific=('xyz','abc',03004021)`

[where xyz and abc are environment variable names]

---

## Setting dump controls for IBM service

**ras\_minorcode\_action= *value***

Determines the default behavior for gathering documentation about system exception minor codes.

**CEEDUMP**

Captures callback and offsets.

**Tip:** It takes time for the system to take CEEDUMPs and this may cause transaction timeouts. For instance, if the WebSphere variable `transaction_defaultTimeout`

is set to 30 seconds, your application transaction may time out because processing a CEEDUMP can take longer than 30 seconds. To prevent this from happening, either:

- Increase the transaction timeout value, or
- Code

```
ras_minorcode_action=NODIAGNOSTICDATA
```

and make sure the

```
ras_trace_minorCodeTraceBacks
```

variable is not specified.

### TRACEBACK

Captures Language Environment and z/OS UNIX traceback data.

### SVCDUMP

Captures an MVS dump (but will not produce a dump in the client).

### NODIAGNOSTICDATA

Specifies that no diagnostic data will be collected, even if CEEDUMP, TRACEBACK, or SVCDUMP processing occurs because of another WebSphere variable setting. For example, if you code both of the following variables, traceback processing occurs but none of the traceback data is collected:

```
ras_minorcode_action=NODIAGNOSTICDATA
```

```
ras_trace_minorCodeTraceBacks=ALL
```

**Default:** NODIAGNOSTICDATA

### Example:

```
ras_minorcode_action=SVCDUMP
```

**ras\_trace\_minorCodeTraceBacks= *value***

Enables traceback of system exception minor codes. Values are:

ALL|all

Enables traceback for all system exception minor codes.

Enables traceback of system exception minor codes. Values are:

- *minor\_code* Enables traceback for a specific minor code.

### Example: Type

```
1234
```

for minor code

```
C9C21234
```

- (*null value*) The default. This setting will not cause gathering of a traceback.

**Default:** (null value)

### Example:

```
ras_trace_minorCodeTraceBacks=all
```

---

## Debugging Service details

Use this page to view and modify the settings used by the Debugging Service.

To view this administrative console page, click **Servers > Application Servers > server > Debugging Service**.

## Startup

Specifies whether the server will attempt to start the Debug service when the server starts.

## JVM debug port

Specifies the port that the Java virtual machine will listen on for debug connections.

## JVM debug arguments

Specifies the debugging argument string used to start the JVM in debug mode.

## Debug class filters

Specifies an array of classes to ignore during debugging. When running in step-by-step mode, the debugger will not stop in classes that match a filter entry.

## BSF debug port

Specifies the port that the BSF Debug Manager listens on.

## BSF logging level

Specifies the level of logging provided by the BSF Debug Manager. The valid range is 0-3, with 3 being the highest level of logging.

---

## Configuration problem settings

Use this page to identify and view problems that exist in the current configuration.

To view this administrative console page, click **Troubleshooting > Configuration Problems** in the console navigation tree.

Click a configuration problem in the Configuration Problems table to see more information about the problem.

## Configuration document validation

Use these fields to specify the level of validation to perform on configuration documents.

## Cross validation

Enables cross validation of configuration documents.

Enabling cross validation enables comparison of configuration documents for conflicting settings.

## Configuration Problems

Displays current configuration problem error messages. Click a message for detailed information about the problem.

## Scope

Sorts the configuration problem list by the configuration file where each error occurs. Click a message for detailed information about the problem.

### Fields that explain each configuration problem:

The following informational fields appear when you click a configuration problem message.

## Message

Displays the message returned from the validator.

## Explanation

A brief explanation of the problem.

## User action

Specifies the recommended action to correct the problem.

## Target Object

Identifies the configuration object where the validation error occurred.

## Severity

Indicates the severity of the configuration error, with 1 being a severe error. Severity decreases as the severity descriptor increases.

## Local URI

Specifies the local URI of the configuration file where the error occurred.

## Full URI

Specifies the full URI of the configuration file where the error occurred.

## Validator classname

The classname of the validator reporting the problem.





---

## Chapter 8. Diagnosing and fixing problems: Resources for learning

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

### The WebSphere Application Server support page

The official site for providing tools and sharing knowledge about WebSphere Application Server problems is the WebSphere Application Server support page: <http://www.ibm.com/software/webservers/appserv/support.html>. Among the features it provides are:

- 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 Application Server technical support personnel.
- A link *All fixes, fix packs, and tools* provides free WebSphere Application Server maintenance upgrades and problem determination tools.
  - *fixes* are software patches which address specific WebSphere Application Server defects. Selecting a specific defect from the list in the *All fixes, fix packs, and tools* page takes you to a description of what problem the fix addresses.
  - Fix packs are bundles of multiple fixes, tested together and released as a maintenance upgrade to WebSphere Application Server. If you select a fix pack from this page, you are taken to a page describing the target platform, WebSphere Application Server prerequisite level, and other related information. Selecting the *list defects* link on that page displays a list of the fixes which the fix pack includes. If you intend to install an fix which is part of a fix pack, it is usually better to upgrade to the complete fix pack rather than to just install the individual fix.
  - Tools are free programs that help you analyze the configuration, behavior and performance of your WebSphere Application Server installation.

### Accessing WebSphere Application Server support page resources

Some resources on the WebSphere Application 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 do not already have an ID. When registering, you are asked for your contract number, which is supplied as part of a WebSphere Application Server purchase.

### WebSphere Developer Domain

The Developer Domains are IBM-supported sites 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. You can reach the WebSphere Developer Domain at <http://www7b.software.ibm.com/wsdd/>.

### **The IBM Support page**

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

---

## Chapter 9. Troubleshooting by task: What are you trying to do?

This section provides troubleshooting information based on the task you were trying to accomplish when the problem occurred. To find more information about your problem, select a task category from the list below.

---

### Troubleshooting installation problems

Select the problem you are having with the WebSphere Application Server installation:

- The “Installation completes, but Sample applications do not work” on page 96.
- I get an “Error when installing WebSphere Application Server for z/OS”
- The “Installation completes but the administrative console does not start”

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

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

#### Error when installing WebSphere Application Server for z/OS

**CEE3250C The system or user abend SCC3 R=00020001 was issued. From compile unit BBODBDLD at entry point BBODBDLD at compile unit offset +xxxxxxx at entryoffset +xxxxxxx at address xxxxxxxx**

This error occurs when SBBOLPA cannot be found in either LPA or STEPLIB. For example, you would see this error if SBBOLPA resides in the link list.

To resolve this problem, place SBBOLPA in LPA or STEPLIB when installing WebSphere Application Server for z/OS and OS/390.

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

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

#### Installation completes but the administrative console does not start

What kind of problem are you having?

- An “Internal Server Error”, “Page cannot be found”, 404, or similar error occurs trying to view the administrative console.
- An “Unable to process login. Check user ID and password and try again. ” error occurs when trying to access console page.
- The directory paths in the console contain strange characters.

If you can bring up the browser page, but the console behavior is inconsistent, error prone, or unresponsive, try upgrading your browser. Older browsers might not support all the features of the administrative console. For a listing of supported

Web browsers, see

<http://www.ibm.com/software/webservers/appserv/doc/latest/prereq.html>.

### **An "Internal Server Error", "Page cannot be found", 404, or similar error occurs trying to view the administrative console**

Here are some steps to try if you are unable to view the administrative console:

- If you are using the deployment manager for a multi-node configuration, run the **startManager** command from the *Network\_Deployment\_install\_dir\bin* directory.
- View the file for the application server or the deployment manager to verify that the server supporting the administrative console has started.
- Verify that the application server which supports the administrative console is up and running.
  - For a base configuration, the administrative console is deployed by default on server1. Before viewing the administrative console, you must verify that server1 is running. Do so by issuing the following command on the MVS console to list active processes:

```
D A,L
```

**Note:** See z/OS MVS System Commands for information on how to use MVS operator commands.

Check for the application server procedure name of BB05ACR with the server short name of BB0S001. If it is not running, enter the following command on the MVS console:

```
START appserver_proc_name,JOBNAME=server_short_name,
 ENV=cell_short_name.node_short_name.server_short_name
```

Example:

```
START BB05ACR,JOBNAME=BBOS001,ENV=PLEX1.SY1.BBOS001
```

- For a network deployment configuration, the deployment manager runs the administrative console. For example, to start the deployment manager, you can issue the following command from the MVS console:

```
START BB05DCR,JOBNAME=BBODMGR,ENV=PLEX1.PLEX1.BBODMGR
```
- View the joblog or sysprint for the application server or deployment manager to verify that the server supporting the administrative console has started.
- Check the URL you use to view the console. By default, it is [http://server\\_name:9080/admin](http://server_name:9080/admin), where *server\_name* is the host name.
- Try to eliminate connection, address and firewall issues by pinging the server machine from a command prompt, using the server name in the URL.

### **An "Unable to process login. Check user ID and password and try again." error occurs when trying to access console page**

This error indicates that security is enabled for WebSphere Application Server, and that the user ID or password supplied is either invalid or not authorized to access the console.

To access the console:

- If you are the administrator, use the ID defined as the security administrative ID. This ID is stored in the WebSphere Application Server file security.xml.

## **Installation completes, but Sample applications do not work**

If the WebSphere Application Server installation program completes successfully, but the Sample applications do not run:

- Browse the application server logs for clues.

- Look up any error or warning messages in the message table by clicking on the **Reference** view of the information center navigation and expanding the **Messages** heading.
- Verify any security-related problems. The error message "MSG0508E: The JMS Server security service was unable to authenticate userid" might display in the error log file when starting an application server.

- **5.1+** You may see the error "com.ibm.mq.jms.BrokerCommandFailedException: Broker command failed: 3008 Reason code 3008" and related errors in the System.out log file after starting the application server which contains the sample applications. This can occur if the userid of the administrative console user is more than 12 characters. The full message will be similar to the following:

```
[9/9/03 9:53:41:234 CST] d05dcb1 ConnectionEve A J2CA0056I: The
Connection Manager received a fatal connection error from the Resource
Adaptor for resource Sample/JMS/TCF. The exception which was received is
com.ibm.mq.jms.BrokerCommandFailedException: Broker command failed: 3008
Reason code 3008 [9/9/03 9:53:41:281 CST] d05dcb1 MDBListenerIm W
WMSG0019E: Unable to start MDB Listener PSSampleMDB, JMSDestination
Sample/JMS/listen : com.ibm.mq.jms.BrokerCommandFailedException: Broker
command failed: 3008 Reason code 3008
```

If this appears to be the problem, try logging out of the administrative console, then log on again with a shorter ID (unless security is enabled, you may enter any alphanumeric value as the user ID). You can also try starting the server from the command line using the command `startserver servername` or `startserver.sh servername`.

---

## Troubleshooting testing and first time run problems

Select the problem you are having with testing or the first run of deployed code for WebSphere Application Server:

- "A web resource does not display."
- "Cannot access a data source" on page 103.
- "Cannot access an enterprise bean from a servlet, a JSP file, a stand-alone program, or another client" on page 120.
- "Cannot look up an object hosted by WebSphere Application Server from a servlet, JSP file, or other client" on page 123.
- "Access problems after enabling security" on page 127.
- "Errors after enabling security" on page 130.
- "Errors after configuring or enabling Secure Sockets Layer" on page 134.
- "Errors in messaging" on page 137.
- "Errors returned to a client sending a SOAP request" on page 138.
- "Errors connecting to WebSphere MQ and creating WebSphere MQ queue connection factory" on page 139.

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

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

### A web resource does not display

If you are not able to display a resource in your browser, follow these steps:

1. Verify that your HTTP server is healthy by accessing the URL `http://server_name` from a browser and seeing whether the Welcome page appears. This action indicates whether the HTTP server is up and running, regardless of the state of WebSphere Application Server.
2. If the HTTP server Welcome page does not appear, that is, if you get a browser message like page cannot be displayed or something similar, try to diagnose your Web server problem.
3. If the HTTP server appears to function, the Application Server might not be serving the target resource. Try accessing the resource directly through the Application Server instead of through the HTTP server.

If you cannot access the resource directly through the Application Server, Verify that the URL used to access the resource is correct.

If the URL is incorrect and it is created as a link from another JSP file, servlet, or HTML file, try correcting it in the browser URL field and reloading, to confirm that the problem is a malformed URL. Correct the URL in the "from" HTML file, servlet or jsp file.

If the URL appears to be correct, but you cannot access the resource directly through the Application Server, verify the health of the hosting Application Server and Web module:

- a. View the hosting Application Server and Web module in the administrative console to verify that they are up and running.
  - b. Copy a simple HTML or JSP file (such as `SimpleJsp.jsp` in the WebSphere Application Server directory structure) to your Web module document root, and try to access it. If successful, the problem is with your resource. View the logs of your Application Server to find out why your resource cannot be found or served
4. If you can access the resource directly through the Application Server, but not through an HTTP server, the problem lies with the HTTP plug-in -- the component that communicates between the HTTP server and the WebSphere Application Server.
  5. If the JSP file and the servlet output are served, but not static resources such as `.html` and image files, see the steps for enabling file serving.
  6. If some kinds of resources display correctly, but you cannot display a servlet by its class name:
    - Verify that the servlet is in a directory in the Web module class path, such as in the `/Web_module_name.war/WEB-INF/classes` directory.
    - Verify that you specify the full class name of the servlet, including its package name, in the URL.
    - Verify that `/servlet` precedes the class name in the URL. For example, if the root context of a Web module is "myapp", and the servlet is `com.mycom.welcomeServlet`, then the URL reads:  
`http://hostname/myapp/servlet/com.mycom.welcomeServlet`
    - For servlets or other resources served by mapped URLs, the URL is `http://hostname/web module context root/mappedURL`.

### Accessing a Web resource through the application server and bypassing the HTTP server

Starting with WebSphere Application Server Version 4.0, you can bypass the HTTP server and access a web resource through the application server. It is not recommended to serve a production Web site in this way, but it provides a good diagnostic tool when it is not clear whether a problem resides in the HTTP server, WebSphere Application Server, or the HTTP plug-in.

To access a Web resource through the Application Server:



1. Determine the port of the HTTP service in the target Application Server.
  - a. In the WebSphere administrative console, click **Servers>Manage Application Servers**.
  - b. Select the target server, then under Additional Properties click **Web Container**.
  - c. Under the Additional Properties of the Web Container, click **HTTP Transports**. You see the ports listed for virtual hosts served by the Application Server.
2. Use the HTTP transport port number of the Application Server to access the resource from a browser. For example, if the port is 9080, the URL is `http://hostname:9080\myAppContext\myJSP.jsp`.
3. If you are still unable to access the resource, verify that the HTTP transport port is in the "Host Alias" list:
  - a. Click **Application Servers > Your\_ApplicationServer > Web Container > HTTP Transports** to check the Default virtual host and the HTTP transport ports used by this Application Server.
  - b. Click **Environment > Manage Virtual Hosts > default host > Host Aliases** to check if the HTTP transport port exists. Add an entry if necessary. For example, if the HTTP port for your application is server is 9080, add a host alias of `*:9082`.

## Errors starting an application

What kind of error do you see when you start an application?

- HTTP server and Application Server are working separately, but requests are not passing from HTTP server to Application Server
- File serving problems
- Graphics do not appear on the JavaServer Pages (JSP) file or servlet output
- SRVE0026E: [Servlet Error]-[Unable to compile class for JSP error on JSP]
- 
- Message similar to "Message: /jspname.jsp(9,0) Include: Mandatory attribute page missing" displays when trying to access JSP file
- 
- The JSP batch compiler fails with the message "Enterprise Application [application name you typed in] not found"
- There is a translation problem with non-English browser input
- Scroll bars do not appear around items in the browser window
- A "Page cannot be displayed... server not found or DNS error" error displays when attempting to browse a Java Server Page (JSP) using Internet Explorer

### HTTP server and Application Server are working separately, but requests are not passing from HTTP server to Application Server

If your HTTP server appears to be functioning correctly, and the Application Server also works on its own, but browser requests sent to the HTTP server for pages are not being served, a problem exists in the WebSphere Application Server plug-in.

In this case:

1. Determine whether the HTTP server is attempting to serve the requested resource itself, rather than forwarding it to the WebSphere Application Server.
  - a. Browse the HTTP server access log (*IHS install root\logs\access.log* for IBM HTTP Server). It might indicate that it could not find the file in its own document root directory.
  - b. browse the plug-in log file as described below.

2. Refresh the *install\_dir/config/plugin-cfg.xml* file that determines which requests sent to the HTTP server are forwarded to the WebSphere Application Server, and to which Application Server. You might need to refresh this file:
  - In the WebSphere Application Server administrative console, expand the Environment tree control.
  - Click **Update WebSphere Plugin**.
  - Stop and restart the HTTP server and retry the Web request.
3. Browse the *install\_dir/logs/http\_plugin.log* file for clues to the problem. Make sure the timestamps with the most recent plug-in information stanza, which is printed out when the plug-in is loaded, correspond to the time the Web server started.
4. Turn on plug-in tracing by setting the `LogLevel` attribute in the *install\_dir/config/plugin-cfg.xml* file to `Trace` and reloading the request. Browse the *install\_dir/logs/http\_plugin.log* file. You should be able to see the plug-in attempting to match the request URI with the various URI definitions for the routes in the *plugin-cfg.xml*. Check which rules the plug-in is not matching against and then figure out if you need to add additional ones. If you just recently installed the application you might need to manually regenerate the plug-in configuration to pick up the new URIs related to the new application.

### File serving problems

If text output appears on your JSP- or servlet-supported Web page, but image files do not:

- Verify that your files are in the right place: the **document root** directory of your Web application WebSphere Application Server follows the J2EE standard, which means that the document root is the *Web\_module\_name.war* directory of your deployed Web application. Typically this directory will be found in the *installation\_root/installedApps/nodename/appname.ear* directory or *installation\_root/installedApps/nodename/appnameNetwork.ear* directory.  
If the files are in a subdirectory of the document root, verify that the reference to the file reflects that. That is, if the *invoices.html* file is stored in Windows directory *Web\_module\_name.war\invoices*, then links from other pages in the Web application to display it should read "*invoices\invoices.html*", not "*invoices.html*".
- Verify that your Web application is configured to enable file serving (in other words, that it is enabled to display static resources like image and *.html* files):
  1. Edit the **fileServingEnabled** property in the deployed Web application *ibm-web-ext.xmi* configuration file, typically found in the *install\_root/config/cells/nodename* or *nodenameNetwork/applications/application\_name/deployments/application\_name/Webmodule name/web-inf* directory.

### Graphics do not appear in the JSP file or servlet output

If text output appears on your JSP- or -servlet-supported Web page, but image files do not:

- Verify that your graphic files are in the right place: the **document root** directory of your Web application WebSphere Application Server Version 5 follows the J2EE standard, which means that the document root is the *Web\_module\_name.war* directory of your deployed Web application. Typically this directory is found in the *installation\_root/installedApps/nodename/appname.ear* directory or *installation\_root/installedApps/nodename/appnameNetwork.ear* directory.

If the graphics files are in a subdirectory of the document root, verify that the reference to the graphic reflects that; for example, if the banner.gif file is stored in Windows directory *Web\_module\_name.war/images*, the tag to display it should read: **<img SRC="images/banner.gif">**, not **<img SRC="banner.gif">**.

- Verify that your Web application is configured to enable file serving (that is, display of static resources like image and .html files).
  1. Edit the **fileServingEnabled** property in the deployed Web application *ibm-web-ext.xml* configuration file, typically found in the *install\_root/config/cells/nodename* or *nodenameNetwork/applications/application name/deployments/application name/Webmodule name/web-inf* directory.
  2. After following the previous steps:
    - In the administrative console, expand the **Environment** tree control .
    - Click **Update WebSphere Plugin**.
    - Stop and restart the HTTP server and retry the Web request.

#### SRVE0026E: [Servlet Error]-[Unable to compile class for JSP file

If this error appears in a browser when trying to access a new or modified .jsp file for the first time, the most likely cause is that the JSP file Java source failed (was incorrect) during the javac compilation phase.

Check the log files for a compiler error message, such as:

```
Duplicate variable declaration: int myInt was int myInt
int myInt = 122;
String myString = "number is 122";
static int myStaticInt=22;
int myInt=121;
 ^
```

Fix the problem in the JSP source file, save the source and request the JSP file again.

If this error occurs when trying to serve a JSP file that was copied from another system where it ran successfully, then there is something different about the new server environment that prevents the JSP file from running. Browse the text of the error for a statement like:

```
Undefined variable or class name: MyClass
```

This error indicates that a supporting class or jar file is not copied to the target server, or is not on the class path. Find the MyClass.class file, and place it on the Web module WEB-INF/classes directory, or place its containing .jar file in the Web module WEB-INF/lib directory.

Verify that the URL used to access the resource is correct by completing the following steps:

- For a JSP file, html file, or image file:  
**http://host\_name/Web\_module\_context\_root/subdir under doc root, if any/filename.ext**. The document root for a Web application is the *application\_name.WAR* directory of the installed application.  
–
- For servlets served by class name, the URL is  
**http://hostname/Web\_module\_context\_root/servlet/packageName.className**.  
–

Correct the URL in the "from" HTML file, servlet or JSP file. An HREF with no leading slash (/) inherits the calling resource context. For example:

- an HREF in `http://[hostname]/myapp/servlet/MyServlet` to "ServletB" resolves to `http://hostname/myapp/servlet/ServletB`
- an HREF in `http://[hostname]/myapp/servlet/MyServlet` to `servlet/ServletB` resolves to `http://hostname/myapp/servlet/servlet/ServletB` (an error)
- an HREF in `http://[hostname]/myapp/servlet/MyServlet` to `/ServletB` resolves to `http://hostname/ServletB` (an error, if ServletB requires the same context root as MyServlet)

### **Message like "Message: /jspname.jsp(9,0) Include: Mandatory attribute page missing" appears when attempting to browse JSP file**

It is probable that the JSP file failed during the translation to Java phase. Specifically, a JSPdirective, in this case an Include statement, was incorrect or referred to a file that could not be found.

To correct this problem, fix the problem in the JSP source, save the source and request the JSP file again.

### **The JSP Batch Compiler fails with the message "Enterprise Application [application name you typed in] not found."**

It is probable that the full enterprise application path and name, starting with the .ear subdirectory that resides in the `install_root\config\cells\node_nameNetwork\applications` directory is expected as an argument to the JspBatchCompiler tool, not just the display name. For example:

- `"JspBatchCompiler -enterpriseapp.name sampleApp.ear/deployments/sampleApp"` is correct, as opposed to
- `"JspBatchCompiler -enterpriseapp.name sampleApp"`, which is incorrect.

### **There is a translation problem with non-English browser input.**

If non-English-character-set browser input cannot be translated after being read by a servlet or JSP file, ensure that the request parameters are encoded according to the expected character set before reading. For example, if the site is Chinese, the target .jsp file should have a line:

```
req.setCharacterEncoding("gb2312");
```

before any `req.getParameter()` calls.

This problem affects servlets and jsp files ported from earlier versions of WebSphere Application Server, which converted characters automatically based upon the locale of the WebSphere Application Server.

### **Scroll bars do not appear around items in the browser window**

In some browsers, tree or list type items that extend beyond their allotted windows do not have scroll bars to permit viewing of the entire list.

To correct this problem, right-click on the browser window and click **Reload** from the menu.

## Error "Page cannot be displayed... server not found or DNS error" appears when attempting to browse a Java Server Page (JSP) using Internet Explorer

This error can occur when an HTTP timeout causes the servant to be brought down and restarted. To correct this problem, increase the ConnectionIOTimeout . value:

1. From the administrative console select System Administration > DeploymentManager > Administration Services > Custom Properties
2. Select ConnectionIOTimeout
3. Increase the ConnectionIOTimeout value
4. Click OK.

## Cannot access a data source

What kind of database are you trying to access?

- Oracle
- DB2
- SQL Server
- Cloudscape
- Sybase
- General data access problems

If the errors described in the previous articles do not match the errors you see:

1. Browse the log files of the application server for clues.
2. Browse the Helper Class property of the data source to verify that it is correct and that it is on the WebSphere Application Server class path. Mysterious errors or behavior might result from a missing or misnamed Helper Class name. If WebSphere Application Server cannot load the specified class, it uses a default helper class that might not function correctly with your database manager.

### General data access problems

- An exception "IllegalConnectionUseException" occurs
- WTRN0062E: An illegal attempt to enlist multiple one phase capable resources has occurred.
- ConnectionWaitTimeoutException.
- com.ibm.websphere.ce.cm.StaleConnectionException: [IBM][CLI Driver] SQL1013N The database alias name or database name "NULL" could not be found. SQLSTATE=42705
- java.sql.SQLException: java.lang.UnsatisfiedLinkError:
- "J2CA0030E: Method enlist caught java.lang.IllegalStateException" wrapped in error "WTRN0063E: An illegal attempt to enlist a one phase capable resource with existing two phase capable resources has occurred" when attempting to execute a transaction.
- java.lang.UnsatisfiedLinkError:xaConnect exception when attempting a database operation
- "J2CA0114W: No container-managed authentication alias found for connection factory or datasource *datasource*" when attempting a database operation
- **5.1+** An error is thrown if you use the ws\_ant command to perform the database customization for Structured Query Language in Java on HP platforms

### IllegalConnectionUseException

This error can occur because a connection obtained from a WAS40DataSource is being used on more than one thread. This usage violates the J2EE 1.3 programming

model, and an exception generates when it is detected on the server. This problem occurs for users accessing a data source through servlets or bean-managed persistence (BMP) enterprise beans.

To confirm this problem, examine the code for connection sharing. Code can inadvertently cause sharing by not following the programming model recommendations, for example by storing a connection in an instance variable in a servlet, which can cause use of the connection on multiple threads at the same time.

**WTRN0062E: An illegal attempt to enlist multiple one phase capable resources has occurred**

This error can occur because:

- An attempt was made to share a single-phase connection, when each **getConnection** method has different connection properties; such as the `AccessIntent`. This attempt causes a non-shareable connection to be created.
- An attempt was made to have more than one unshareable connection participate in a global transaction, when the data source is not an XA resource.
- An attempt was made to have a one-phase resource participate in a global transaction while an XA resource or another one-phase resource already participated in this global transaction.
  - Within the scope of a global transaction you try to get a connection more than once and at least one of the resource-refs you use specifies that the connection is unshareable, and the data source is not configured to support two-phase commit transactions. It does not support an `XAResource`. If you do not use a resource-ref, you default to unshareable connections.
  - Within the scope of a global transaction you try to get a connection more than once and at least one of the resource-refs you use specifies that the connection is shareable and the data source is not configured to support two-phase commit transactions. That is, it does not support an `XAResource`. In addition, even though you specify that connections are shareable, each `getConnection` request is made with different connection properties (such as `IsolationLevel` or `AccessIntent`). In this case, the connections are not shareable, and multiple connections are handed back.
  - Multiple components (servlets, session beans, BMP entity beans, or CMP entity beans) are accessed within a global transaction. All use the same data source, all specify shareable connections on their resource-refs, and you expect them to all share the same connection. If the properties are different, you get multiple connections. `AccessIntent` settings on CMP beans change their properties. To share a connection, the `AccessIntent` setting must be the same. For more information about CMP beans sharing a connection with non-CMP components, see the *Data access application programming interface support* and *Example: Accessing data using IBM extended APIs to share connections between container-managed and bean-managed persistence beans* topics in the `DataAccess` section of the information center.

To correct this error:

- Check what your client code passes in with its `getConnection` requests, to ensure they are consistent with each other.
- Check the JDBC provider implementation class from the Manage JDBC resource panel of the administrative console to ensure that it is a class that supports XA-type transactions.



## ConnectionWaitTimeoutException accessing a data source or resource adapter

If your application receives exceptions like a `com.ibm.websphere.ce.cm.ConnectionWaitTimeoutException` or `com.ibm.websphere.ce.j2c.ConnectionWaitTimeoutException` when attempting to access a WebSphere Application Server data source or JCA-compliant resource adapter, respectively, some possible causes are:

- The maximum number of connections for a given pool is set too low. The demand for concurrent use of connections is greater than the configured maximum value for the connection pool. One indication that this situation is the problem is that you receive these exceptions regularly, but your CPU utilization is not high. This exception indicates that there are too few connections available to keep the threads in the server busy.
- Connection Wait Time is set too low. Current demand for connections is high enough such that sometimes there is not an available connection for short periods of time. If your connection wait timeout value is too low, you might timeout shortly before a user returns a connection back to the pool. Adjusting the connection wait time can give you some relief. One indication of this problem is that you use close to the maximum number of connections for an extended period and receiving this error regularly.
- You are not closing some connections or you are returning connections back to the pool at a very slow rate. This situation can happen when using unshareable connections, when you forget to close them, or you close them long after you are finished using them, keeping the connection from returning to the pool for reuse. The pool soon becomes empty and all applications get `ConnectionWaitTimeoutExceptions`. One indication of this problem is you run out of connections in the connection pool and you receive this error on most requests.
- You are driving more load than the server or backend system have resources to handle. In this case you must determine which resources you need more of and upgrade configurations or hardware to address the need. One indication of this problem is that the application or database server CPU is nearly 100% busy.

To correct these problems, either:

- Modify an application to use fewer connections
- Properly close the connections.
- Change the pool settings of `MaxConnections` or `ConnectionWaitTimeout`.
- Adjust resources and their configurations.

**com.ibm.websphere.ce.cm.StaleConnectionException: [IBM][CLI Driver] SQL1013N The database alias name or database name "NULL" could not be found. SQLSTATE=42705**

This error occurs when a data source is defined but the `databaseName` attribute and the corresponding value are not added to the custom properties panel.

To add the `databaseName` property:

1. Click **Resources>Manage JDBC Providers** link in the administrative console.
2. Select the JDBC provider that supports the problem data source.
3. Select **Data Sources** and then select the problem data source.
4. Under **Additional properties** click **Custom Properties**.
5. Select the `databaseName` property, or add one if it does not exist, and enter the actual database name as the value.
6. Click **Apply** or **OK**, and then click **Save** from the action bar.
7. Access the data source again.



### **java.sql.SQLException: java.lang.UnsatisfiedLinkError:**

This error indicates that the directory containing the binary libraries which support a database are not included in the LIBPATH environment variable for the environment in which the WebSphere Application Server starts.

The path containing the DBM vendor libraries vary by dbm. One way to find them is by scanning for the missing library specified in the error message. Then you can correct the LIBPATH variable to include the missing directory, either in the .profile of the account from which WebSphere Application Server is executed, or by adding a statement in a .sh file which then executes the startServer program.

**"J2CA0030E: Method enlist caught java.lang.IllegalStateException" wrapped in error "WTRN0063E: An illegal attempt to enlist a one phase capable resource with existing two phase capable resources has occurred" when attempting to execute a transaction.**

This error can occur when last participant support (LPS) is missing or disabled. LPS allows a one-phase capable resource and a two-phase capable resource to enlist within the same transaction.

LPS is only available if the following are true:

- WebSphere Application Server Programming Model Extensions (PME) is installed. PME is included in the Application Server Integration Server product.
- The Additional Integration Server Extensions option is enabled when PME is installed. If you perform a typical installation, this function is enabled by default. If you perform a custom installation, you have the option to disable this function, which disables LPS.

### **java.lang.UnsatisfiedLinkError:xaConnect exception when attempting a database operation**

This problem has two main causes:

- The most common cause is that the jdbc driver which supports connectivity to the database is missing, or is not the correct version, or that native libraries which support the driver are on the system's path.
  - To resolve this problem on a Windows platform, verify that the JDBC driver jar file is on the system PATH environment variable:
    - If you are using DB2, verify that at least the DB2 client product has been installed on the WebSphere host
      - On DB2 version 7.2 or earlier, the file where the client product is installed on the WebSphere Application Server is db2java.zip. Verify that the usejdbc2.bat program has been executed after the database install and after any upgrade to the database product.
      - On DB2 version 8.1 or later, use the DB2 Universal JDBC Provider Driver when defining a JDBC provider in WebSphere Application Server. The driver file is db2jcc.jar. If you use the type 2 (default) option, verify that at least the DB2 client product is installed on the WebSphere Application Server host. If you specify the type 4 option, the DB2 client does not need to be installed, but the file db2jcc.jar still must be present.

When specifying the location of the driver file, it is recommended to that you specify the path and file name of the target DB2 installation, rather than simply copying the file to a local directory, if possible. Otherwise, you may be exposed to problems if the target DB2 installation is

upgraded and the driver used by WebSphere Application Server is not. If you choose **DB2 Legacy CLI-based type 2 JDBC Driver** when defining a JDBC provider in WebSphere Application Server, then you must still follow the steps to ensure that you have the correct version of the db2java.zip file (see instructions for DB2 7.2 or earlier).

- On a Unix platform, ensure that any native libraries required to support the database client of your database product are specified in the LD\_LIBRARY\_PATH environment variable in the profile of the account under which WebSphere Application Server executes.

If you are using DB2 The native library is libdb2jdbc.so. The best way to ensure that this library is accessed correctly by WebSphere is to call the db2profile script supplied with DB2 from the .profile script of the account (such as "root") under which WebSphere runs.

- If you are using DB2 version 7.2 or earlier, ensure that the usejdbc2,script provided with DB2 is called from the profile of the account under which WebSphere Application server is launched.
- If you are using DB2 version 8.1 or later, see the previous instructions for the Windows operating system.
- If the database manager is DB2, you may have chosen the option to create a 64-bit instance. A 64-bit configuration is not supported. If this has happened, remove the database instance and create a new one with the default 32-bit setting.

**Note:** For general help in configuring JDBC drivers and data sources in WebSphere Application Server, see the topic Accessing data from applications.

### **"J2CA0114W: No container-managed authentication alias found for connection factory or datasource *datasource*" when attempting a database operation**

This error might occur in the SystemOut.log file when you run an application to access a data source after creating the data source using JACL script.

The error message occurs because the JACL script did not set container-managed authentication alias for CMP connection factory. The JACL is missing the following line:

```
$AdminConfig create MappingModule $cmpConnectorFactory
"{mappingConfigAlias DefaultPrincipalMapping} {authDataAlias $authDataAlias}
```

To correct this problem, add the missing line to the JACL script and run the script again. See Example: Creating a JDBC provider and data source using Java Management Extensions API and the scripting tool for a sample JACL script.

### **An error is thrown if you use the ws\_ant command to perform the database customization for Structured Query Language in Java on HP platforms **5.1 +****

If you use the ws\_ant command to perform the database customization for Structured Query Language in Java (SQLJ) on HP platforms, you can receive an error similar to the following:

```
[java] [ibm][db2][jcc][sqlj]
[java] [ibm][db2][jcc][sqlj] Begin Customization
[java] [ibm][db2][jcc][sqlj] encoding not supported!!
```

The cause of this error might be that your databases were created using the HP default character set. The Java Common Client (JCC) driver depends on the software development kit (SDK) to perform the codepage conversions. The SDK shipped with this product, however, does not support the HP default codepage.

You need to set your LANG to the ISO locale before creating the databases. It should be similar to the following:

```
export LANG=en_US.iso88591
```

Refer to the DB2 Tech Notes at <http://www-3.ibm.com/software/data/db2/udb/ad/v8/bldg/t0004877.htm> for details.

## Problems accessing an Oracle data source

**What kind of error do you see when you try to access your Oracle-based data source**

- An Invalid Oracle URL is specified
- "DSRA0080E: An exception was received by the Data Store Adapter. See original exception message: ORA-00600" when connecting to or using an Oracle data source.
- "DSRA8100E: Unable to get a {0} from the DataSource. Explanation: See the linkedException for more information."
- An "Error while trying to retrieve text for error" message occurs when connecting to an Oracle data source.
- A java.lang.UnsatisfiedLinkError occurs when connecting to an Oracle data source.
- The exception java.lang.NullPointerException or an "internal error: oracle.jdbc.oci8.OCIEnv" occurs when connecting to an Oracle data source.
- WSVR0016W: Classpath entry, \${ORACLE\_JDBC\_DRIVER\_PATH}/classes12.zip, in Resource, Oracle JDBC Thin Driver, located at cells/BaseApplicationServerCell/nodes/wasrtp/resources.xml has an invalid variable.

### An invalid Oracle URL is specified

This error might be caused by an incorrectly specified URL on the URL property of the target data source.

Examine the URL property for the data source object in the administrative console. For the 8i OCI driver, verify that **oci8** is used in the URL. For the 9i OCI driver, you can use either **oci8** or **oci**.

Examples of Oracle URLs:

- For the thin driver: jdbc:oracle:thin:@hostname.rchland.ibm.com:1521:IBM
- For the thick (OCI) driver: jdbc:oracle:oci8:@tnsname1

**"DSRA0080E: An exception was received by the data store adapter. See original exception message: ORA-00600" when connecting to or using an Oracle data source "DSRA0080E: An exception was received by the data store adapter. See original exception message: ORA-00600" when connecting to or using an Oracle data source "DSRA0080E: An exception was received by the Data Store Adapter. See original exception message: ORA-00600" when connecting to or using an Oracle data source**

A possible reason for this exception is that the version of the Oracle JDBC driver being used is older than the Oracle database. It is possible that more than one version of the Oracle JDBC driver is configured on the WebSphere Application Server.

Examine the version of the JDBC driver. Sometimes you can determine the version by looking at the class path to determine what directory the driver is in.

If you cannot determine the version this way, use the following program to determine the version. Before running the program, set the class path to the location of your JDBC driver files.

```
import java.sql.*;
import oracle.jdbc.driver.*;
class JDBCVersion
{
 public static void main (String args[])
 throws SQLException
 {
 // Load the Oracle JDBC driver
 DriverManager.registerDriver(new oracle.jdbc.driver.OracleDriver());
 // Get a connection to a database
 Connection conn = DriverManager.getConnection
 ("jdbc:oracle:thin:@appaloosa:1521:app1","sys","change_on_install");
 // Create Oracle DatabaseMetaData object
 DatabaseMetaData meta = conn.getMetaData();
 // gets driver info:
 System.out.println("JDBC driver version is " + meta.getDriverVersion());
 }
}
```

If the driver and the database are at different versions, replace the JDBC driver with the correct version. If multiple drivers are configured, remove any that occur at the incorrect level.

**DSRA8100E: Unable to get a {0} from the DataSource. Explanation: See the [linkedException](#) for more information.**

When using an oracle thin driver, Oracle throws a "java.sql.SQLException: invalid arguments in call" error if no user name or password is specified when getting a connection. If you see this error while running WebSphere Application Server, the alias is not set.

To remove the exception, define the alias on the data source.

**"Error while trying to retrieve text for error" error when connecting to an Oracle data source**

The most likely cause of this error is that the Oracle 8i OCI driver is being used with an ORACLE\_HOME property that is either not set or is set incorrectly.

To correct the error, examine the user profile that WebSphere Application Server is running under to verify that the \$ORACLE\_HOME environment variable is set correctly.

### **"java.lang.UnsatisfiedLinkError:" connecting to an Oracle data source**

The environment variable LIBPATH might not be set or is set incorrectly, if your data source throws an **UnsatisfiedLinkError** error, and the full exception indicates that the problem is related to an Oracle module, as in the following examples.

Example of invalid an LIBPATH for the 8i driver:

```
Exception in thread "main" java.lang.UnsatisfiedLinkError:
/usr/WebSphere/AppServer/java/jre/bin/libocijdbc8.so:
load ENOENT on shared library(s)
/usr/WebSphere/AppServer/java/jre/bin/libocijdbc8.so libclntsh.a
```

Example of an invalid LIBPATH for the 9i driver:

```
Exception in thread "main" java.lang.UnsatisfiedLinkError:
no ocijdbc9 (libocijdbc9.a or .so) in java.library.path
at java.lang.ClassLoader.loadLibrary(ClassLoader.java:Compiled Code)
at java.lang.Runtime.loadLibrary0(Runtime.java:780)
```

To correct the problem, examine the user profile under which the WebSphere Application Server is running to verify that the LIBPATH environment variable includes Oracle libraries. Scan for the libocijdbc8.so file to find the right directory.

### **java.lang.NullPointerException referencing 8i classes, or " internal error: oracle.jdbc.oci8. OCIEnv" connecting to an Oracle data source**

The problem might be that the 9i OCI driver is being used on an AIX 32-bit machine, the LIBPATH is set correctly, but the ORACLE\_HOME environment variable is not set or is set incorrectly. You can encounter an exception similar to either of the following when your application attempts to connect to an Oracle data source:

Exception example for the java.lang.NullPointerException:

```
Exception in thread "main" java.lang.NullPointerException
at oracle.jdbc.oci8.OCIEnv.check_error(OCIEnv.java:1743)
at oracle.jdbc.oci8.OCIEnv.getEnvHandle(OCIEnv.java:69)
at oracle.jdbc.oci8.OCIEnv.logon(OCIEnv.java:452)
at oracle.jdbc.driver.OracleConnection.<init>(OracleConnection.java:287)
```

Exception example for the java.sql.SQLException:

```
Exception in thread "main" java.sql.SQLException:
internal error: oracle.jdbc.oci8. OCIEnv@568b1d21
at oracle.jdbc.dbaccess.DBError.throwSQLException(DBError.java:184)
at oracle.jdbc.dbaccess.DBError.throwSQLException(DBError.java:226)
at oracle.jdbc.oci8.OCIEnv.getEnvHandle(OCIEnv.java:79)
```

To correct the problem, examine the user profile that WebSphere Application Server is running under to verify that it has the \$ORACLE\_HOME environment variable set correctly, and that the \$LIBPATH includes \$ORACLE\_HOME/lib.

**WSVR0016W: Classpath entry, \${ORACLE\_JDBC\_DRIVER\_PATH}/classes12.zip, in Resource, Oracle JDBC Thin Driver, located at cells/BaseApplicationServerCell/nodes/wasrtp/resources.xml has an invalid variable**

This error occurs when there is no environment variable defined for the property, ORACLE\_JDBC\_DRIVER\_PATH.

Verify this problem in the administrative console. Go to **Environment > Manage WebSphere Variables** to verify whether the variable ORACLE\_JDBC\_DRIVER\_PATH is defined.

To correct the problem, click **New** and define the variable. For example, name : **ORACLE\_JDBC\_DRIVER\_PATH** , value : **c:\oracle\jdbc\lib** Use a value that names the directory in your operating system and directory structure that contains the classes12.zip file.

## Problems accessing a DB2 database

**What kind of problem are you having accessing your DB2 database?**

- SQL0567N "DB2ADMIN" is not a valid authorization ID. SQLSTATE=42602.
- SQL0805N Package "*package name*" was not found.
- SQL0805N Package "NULLID.SQLLC300" was not found. SQLSTATE=51002.
- SQLException, with ErrorCode -99,999 and SQLState 58004, with Java "StaleConnectionException: COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver] CLI0119E Unexpected system failure. SQLSTATE=58004" using WAS40-type data source.
- 
- CLI0119E System error. SQLSTATE=58004 - DSRA8100 : Unable to get a XAconnection, or DSRA0011E: Exception: COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver] CLI0119E Unexpected system failure. SQLSTATE=58004.
- COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver][DB2/NT] SQL0911N The current transaction has been rolled back because of a deadlock or timeout. Reason code "2". SQLSTATE=40001.
- "COM.ibm.db2.jdbc.DB2ConnectionPoolDataSource" could not be found for data source (*data\_source*).
- Error 10500 (E\_Fatal) returns when an attempt is made to either publish or inquire on an entity exceeding 255 characters in one of its fields.
- **5.1+** Database access does not work after initial creation of JDBC provider (DB2 for z/OS JDBC Provider or DB2 for z/OS JDBC Provider (XA)).
- **5.1+** A NullPointerException error is thrown if the serverName property is not provided when running type 4 DB2 Universal JDBC driver

**SQL0567N "DB2ADMIN " is not a valid authorization ID. SQLSTATE=42602**

If you encounter this error when attempting to access a DB2 Universal Database (UDB):

1. Verify that your user name and password in the data source properties page in the administrative console are correct.
2. Ensure that the user ID and password do not contain blank characters before, in between, or after.

**SQL0805N Package *package-name* was not found**

Possible reasons for these exceptions:



- If the package name is NULLID.SQLLC300, see SQL0805N Package "NULLID.SQLLC300" was not found. SQLSTATE=51002. for the reason.
- You are attempting to use an XA-enabled JDBC driver on a DB2 database that is not XA-ready.

To correct the problem on a DB2 Universal Database (UDB), run this one-time procedure, using the db2cmd interface while connected to the database in question:

1. **DB2 bind @db2ubind.lst blocking all grant public**
2. **DB2 bind @db2cli.lst blocking all grant public**

The db2ubind.lst and db2cli.lst files are in the bnd directory of your DB2 installation root. Run the commands from that directory.

### **SQL0805N Package "NULLID.SQLLC300" was not found. SQLSTATE=51002**

This error can occur because:

- The underlying database was dropped and recreated.
- DB2 was upgraded and its packages are not rebound correctly.

To resolve this problem, rebind the DB2 packages by running the db2cli.lst script found in the bnd directory. For example: db2>@db2cli.lst.

### **SQLException, with ErrorCode -99,999 and SQLState 58004, with Java "StaleConnectionException: COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver] CLI0119E Unexpected system failure. SQLSTATE=58004", when using WAS40-type data source**

An unexpected system failure usually occurs when running in XA mode (two-phase commit). Among the many possible causes are:

- An invalid username or password was provided.
- The database name is incorrect.
- Some DB2 packages are corrupted.

To determine whether you have a user name or password problem, look in the db2diag.log file to view the actual error message and SQL code. A message like the following example, with an SQLCODE of -1403, indicates an invalid user ID or password:

```
2002-07-26-14.19.32.762905 Instance:db2inst1 Node:000
PID:9086(java) Appid:*LOCAL.db2inst1.020726191932
XA DTP Support sqlxa_open Probe:101
DIA4701E Database "POLICY2" could not be opened
for distributed transaction processing.
String Title: XA Interface SQLCA PID:9086 Node:000
SQLCODE = -1403
```

To resolve these problems:

1. Correct your user name and password. If you specify your password on the GUI for the data source, ensure that the user name and password you specify on the bean are correct. The user name and password you specify on the bean overwrite whatever you specify when creating the data source.
2. Use the correct database name.
3. Rebind the packages (in the bnd directory) as follows:
 

```
db2connect to dbname
c:\SQLLIB\bnd>DB2 bind @db2ubind.lst blocking all grant public
c:\SQLLIB\bnd>DB2 bind @db2cli.lst blocking all grant public
```
4. Ensure that the \WebSphere\AppServer\properties\wsj2cdpm.properties file has the right user ID and password.



**CLI0119E System error. SQLSTATE=58004 - DSRA8100 : Unable to get a XAconnection or DSRA0011E: Exception: COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver] CLI0119E Unexpected system failure. SQLSTATE=5800**

If you encounter this error when attempting to access a DB2 Universal Database (UDB) data source:

1. Check your user name and password custom properties in the data source properties page in the administrative console. Verify that they are correct.
2. Ensure the user ID and password do not contain any blank characters, before, in between, or after.
3. Check that the `WAS.policy` file exists for the application.
4. View the entire exception listing for an underlying SQL error, and look it up using the DBM vendor message reference.

If you encounter this error while running DB2 on Red Hat Linux, the **max queues system wide** parameter is too low to support DB2 while it acquires the necessary resources to complete the transaction. When this problem exists, the exceptions J2CA0046E and DSRA0010E can precede the exception DSRA8100E.

To correct this problem, edit the `/proc/sys/kernal/msgmni` file to increase the value of the **max queues system wide** parameter to a value greater than 128.

**COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver][DB2/NT] SQL0911N The current transaction has been rolled back because of a deadlock or timeout. Reason code "2". SQLSTATE=40001**

This problem is probably an application-caused DB2 deadlock, particularly if you see an error similar to the following when accessing a DB2 data source:

```
ERROR CODE: -911
COM.ibm.db2.jdbc.DB2Exception: [IBM][CLI Driver][DB2/NT] SQL0911N
The current transaction has been rolled back because of a deadlock or timeout.
Reason code "2". SQLSTATE=40001
```

To diagnose the problem:

1. Execute these DB2 commands:
  - a. `db2 update monitor switches using LOCK ON`
  - b. `db2 get snapshot for LOCKS on dbName >`

The `directory_name\lock_snapshot.log` now has the DB2 lock information.
2. Turn off the lock monitor by executing: `db2 update monitor switches using LOCK OFF`

To verify that you have a deadlock:

1. Look for an application handle that has a lock-wait status, and then look for the ID of the agent holding lock to verify the ID of the agent.
2. Go to that handle to verify it has a lock-wait status, and the ID of the agent holding the lock for it. If it is the same agent ID as the previous one, then you know that you have a circular lock (deadlock).

To resolve the problem:

1. Examine your application and use a less restrictive isolation level if no concurrency access is needed.
2. Use caution when changing the **accessIntent** value to move to a lower isolation level. This change can result in data integrity problems.
3. For DB2/UDB Version 7.2 and earlier releases, you can set the `DB2_RR_TO_RS` flag from the DB2 command line window to eliminate unnecessary deadlocks,

such as when the accessIntent defined on the bean method is too restrictive, for example, PessimisticUpdate. The DB@\_RR\_TO\_RS setting has two impacts:

- If RR is your chosen isolation level, it is effectively downgraded to RS.
- If you choose another isolation level, and the DB2\_RR\_TO\_RS setting is on, a scan skips over rows that are deleted but not committed, even though the row might qualify for the scan. The skipping behavior affects the RR, Read Stability (RS), and Cursor Stability (CS) isolation levels.

For example, consider the scenario where transaction A deletes the row with column1=10 and transaction B does a scan where column1>8 and column1<12. With DB2\_RR\_TO\_RS off, transaction B waits for transaction A to commit or rollback. If transaction A rolls back, the row with column1=10 is included in the result set of the transaction B query. With DB2\_RR\_TO\_RS on, transaction B does not wait for transaction A to commit or rollback. Transaction B immediately receives query results that do not include the deleted row. Setting DB2\_RR\_TO\_RS effectively changes locking behavior, thus avoiding deadlocks.

#### **"COM.ibm.db2.jdbc.DB2ConnectionPoolDataSource" could not be found for data source ({data-source-name})"**

This error usually occurs when the class path of the DB2 JDBC driver is set correctly to \${DB2\_JDBC\_DRIVER\_PATH}/db2java.zip but the environment variable DB2\_JDBC\_DRIVER\_PATH is not set.

To confirm this problem:

1. Go to the **Manage WebSphere Variables** panel.
2. Select **Environment** to verify that there is no entry for the variable DB2\_JDBC\_DRIVER\_PATH.

To correct this problem: Add the variable DB2\_JDBC\_DRIVER\_PATH with **value** equal to the directory path containing the db2java.zip file.

#### **Error 10500 (E\_Fatal) returns when an attempt is made to either publish or inquire on an entity exceeding 255 characters in one of its fields.**

This error usually occurs when you make an attempt to either publish or inquire on an entity exceeding 255 characters in one of its fields. This is less obvious when using non-English characters, as the real limit is reached before 255 visual characters are displayed.

To correct this problem: Accept this as a limitation in using DB2 Version 7 on z/OS. Do not exceed the 255 character limit.

#### **Database access does not work after initial creation of JDBC provider (DB2 for z/OS JDBC Provider or DB2 for z/OS JDBC Provider (XA)). **5.1 +****

When you use the administrative console to create a DB2 z/OS Java Database Connectivity (JDBC) provider or a DB2 z/OS JDBC provider (XA), you will not be able to access the database unless you take the following steps:

- Select the normal DB2 call level interface (CLI)-based JDBC provider option. It uses the same implementation class as both of the z/OS JDBC providers.
- When you create a data source under this provider, change the data store helper value from *DB2DataStoreHelper* to *DB2390DataStoreHelper*.

A `NullPointerException` error is thrown if the `serverName` property is not provided when running type 4 DB2 Universal JDBC driver **5.1+**

|                             |                                                                                                                                                                                                                                                      |
|-----------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Problem</b>              | Instead of an exception with a meaningful message, DB2 is throwing a <code>NullPointerException</code> error.                                                                                                                                        |
| <b>Recommended response</b> | If you are using the Universal JDBC Type 4 driver, verify that you have provided a value for the <code>serverName</code> property. Setting the <code>driverType</code> property to 4 requires that the <code>serverName</code> property be provided. |

For more information about configuring the DB2 Universal JDBC Driver Provider, refer to the DB2 section of the topic *Vendor-specific data sources minimum required settings*.

### Problems accessing a SQL server data source

What kind of problem are you having accessing your SQL Server database?

- ERROR CODE: 20001 and SQL STATE: HY000.
- Application fails with message stating "Cannot find stored procedure..."
- ERROR CODE: SQL5042 when running a Java application

#### ERROR CODE: 20001 and SQL STATE: HY000 accessing SQLServer database

The problem might be that the distributed transaction coordinator service is not started. Look for an error similar to the following example when attempting to access an SQL server database:

```
ERROR CODE: 20001
SQL STATE: HY000
java.sql.SQLException: [Microsoft][SQLServer JDBC Driver]
[SQLServer]xa_open (0) returns -3
at com.microsoft.jdbc.base.BaseExceptions.createException(Unknown Source) ...
at com.microsoft.jdbcx.sqlserver.SQLServerDataSource.getXAConnection
(Unknown Source) ...
```

To confirm this problem:

1. Go to the Windows **Control Panel** and click **Services**(or click **Control Panel > Administrative Tools > Services**)
2. Verify whether the service **Distributed Transaction Coordinator** or **DTC** is started.
3. If not, start the Distributed Transaction Coordinator service.

#### Application fails with message stating "Cannot find stored procedure..." accessing an SQLServer database

This error can occur because the stored procedures for the Java Transaction API (JTA) feature are not installed on the Microsoft SQL Server.

To resolve the problem: Repeat the installation for the stored procedures for the JTA feature, according to the *ConnectJDBC installation guide*.

#### ERROR CODE: SQL5042 when running a Java application

This error can occur when you configure your application to run in the following manner:

1. you use a type 2 (application) driver running on the gateway to the OS 390
2. your application is an XA application.

OS 390 does not use XA, but uses SPM. To resolve the problem:

1. Check your dbm cfg to see that the SPM is not started on the gateway.
2. Assign a port and set the *db2comm* variable to **TCPIP**.
3. Update the dbm cfg value *SPM\_NAME* to use your machine name.
4. Start the SPM on the gateway.

## **Problems accessing a Cloudscape database**

### **What kind of problem are you having accessing your Cloudscape database?**

- Unexpected IO exception wrapped in SQL exception, accessing Cloudscape database.
- The "Select for update" operation on one row causes table to become locked, triggering a deadlock condition.
- "ERROR XSDB6: Another instance of Cloudscape might have already booted the database *databaseName*." error starting application server.
- Error "The version of the IBM Universal JDBC driver in use is not licensed for connectivity to Cloudscape databases"
- Running an application causes a runtime exception which produces an unreadable message.
- Unexpected SQL exception after initially configuring Cloudscape Network Server for use with WebSphere Application Server, version 5.0.2 or later.
- The administrative console displays DB2 settings instead of Cloudscape settings when you create a new Cloudscape JDBC Provider Network Server using the universal JDBC driver provider.

**Tip:** Cloudscape errorCodes (2000, 3000, 4000) indicate levels of severity, not specific error conditions. In diagnosing Cloudscape problems, pay attention to the given sqlState value.

### **Unexpected IOException wrapped in SQLException, accessing Cloudscape database**

This problem can occur because Cloudscape databases use a large number of files. Some operating systems, such as the Solaris Operating Environment, limit the number of files an application can open at one time. If the default is a low number, such as 64, you can get this exception.

If your operating system lets you configure the number of file descriptors, you can correct the problem by setting the number to a high value, such as 1024.

### **The "select for update" operation causes table lock and deadlock when accessing Cloudscape**

If a select for update operation on one row locks the entire table, which creates a deadlock condition, there might be undefined indexes on that table. The lack of an index on the columns you use in the where clause can cause Cloudscape to create a table lock rather than a row level lock.

To resolve this problem, create an index on the affected table.

### **ERROR XSDB6: Another instance of Cloudscape may have already booted the database "database"**

This problem occurs because Cloudscape embedded framework only allows one Java virtual machine (JVM) to access the database instance at a time.

To resolve this problem:

1. Verify that you do not have other JDBC client programs, such as **ij** or **cvview** running on that database instance, when WebSphere Application Server is running.
2. Verify that you do not use the same instance of the database for more than one data source or use the networkServer framework, which doesn't have this limitation.
3. If there are no connections to Cloudscape, delete the db.lck lock file. This file can be found in the directory where the Cloudscape database is mounted, under the schema directory. For example, if the database is mounted at /myCloudscapeDB, issue the command: `rm /myCloudscapeDB/schemaName/db.lck`

Error "The version of the IBM Universal JDBC driver in use is not licensed for connectivity to Cloudscape databases"

### **Error "The version of the IBM Universal JDBC driver in use is not licensed for connectivity to Cloudscape databases"**

At the client runtime, an error similar to the following occurs:

```
The version of the IBM Universal JDBC driver in use is not
licensed for connectivity to Cloudscape databases. To connect
to this DB2 server, please obtain a licensed copy of the IBM DB2
Universal Driver for JDBC and SQLJ. An appropriate license file
db2jcc_license_*.jar for this target platform must be installed to
the application classpath. Connectivity to Cloudscape databases is
enabled by any of the following license files:
{ db2jcc_license_c.jar, b2jcc_license_cu.jar, db2jcc_license_cisuz.jar }
```

The problem occurs because an incorrect JDBC driver jar file name is specified in the class path for JDBC provider. For example, the jar file name may have an extra '\_', as follows:

```
$(UNIVERSAL_JDBC_DRIVER_PATH)/db2jcc_license__cu.jar
```

To resolve the problem:

1. Correct the UNIVERSAL\_JDBC\_DRIVER\_PATH jar file name in the JACL script
2. Restart the cluster.
3. Rerun the client.

### **Running an application causes a runtime exception which produces an unreadable message.**

At client runtime, you may receive a message similar to the following: Caused by: com.ibm.db2.jcc.a.SqlException: DB2 SQL error: SQLCODE: -1, SQLSTATE: 42X05, SQLERRMC: ANNUITYHOLDER20^T42X05

The problem occurs because the property *retrieveMessagesfromServerOnGetMessage*, which is required by WebSphere Application Server, has not been set.

To resolve the problem, on the administrative console

1. Click **Resources** -> **JDBC Providers**
2. Click on a Cloudscape provider
3. Scroll down and click on **Data Sources**
4. Select your data source (or add a new one)
5. Scroll down and select **Custom Properties**
6. If the property *retrieveMessagesFromServerOnGetMessage* already exists, set its value to true. If the property does not exist, select **New** and add the property *retrieveMessagesFromServerOnGetMessage* with a value **true**
7. Rerun the client

The SystemOut.log will now generate readable messages so that you can resolve the underlying problem.

### **Unexpected SQL exception after initially configuring Cloudscape Network Server for use with WebSphere Application Server, version 5.0.2 or later**

#### **Problem**

When running WebSphere Application Server, version 5.0.2 or later, with Cloudscape using the Network Server framework, you receive the following exception:

```
com.ibm.db2.jcc.b.SQLException:
Unexpected throwable caught
java.lang.UnsatisfiedLinkError: DBConnect
```

#### **Recommended response**

Verify that you have supplied a value for the *driverType* data source custom property, and that the value is **not 2**. For more information about configuring Cloudscape Network Server using the Universal JDBC driver, refer to the Cloudscape section of the topic *Vendor-specific data sources minimum required settings*.

### **The administrative console displays DB2 settings instead of Cloudscape settings when you create a new Cloudscape JDBC Provider Network Server using the universal JDBC driver provider.**

#### **Problem**

When creating a new Cloudscape JDBC Provider Network Server using the universal JDBC driver provider through the administrative console in WebSphere Application Server, version 5.0.2 or later, be aware that the administrative console fills in the DB2 universal data sources settings by default instead of the Cloudscape Network Server settings. These erroneous settings include:

- the *DataStoreHelper* class
- the data source custom properties

The provided implementation class, however, is correct. The implementation class for the DB2 universal driver provider happens to be the same as that required for the Cloudscape JDBC Provider Network Server.

## Recommended response

If you want to create the Cloudscape JDBC Provider Network Server using the universal JDBC driver provider, you must modify the entries in the administrative console as follows:

1. Change the `DataStoreHelper` class to:  
`com.ibm.websphere.rsadapter.CloudscapeNetworkServerDataStoreHelper`
2. Supply the required custom properties with these values:
  - `driverType`: 4 (Type 4 driver is the only valid value for Network Server)
  - `serverName`: *your server TCP/IP address or name*
  - `databaseName`: *your database name* (For example, `c:\mydb`)
  - `retrieveMessagesFromServerOnGetMessage`: `true` (This is required by WebSphere Application Server, not the database. Do not use the database default setting, which is `false`.)

You have the option of changing the `portNumber` property, which is the port number on which Network Server listens for connection requests. The default is port 1527. You must either supply another number, or accept the default.

The following data source custom properties are optional:

- `logWriter`
- `traceLevel`
- `traceFile`
- `traceFileAppend`
- `deferPrepares` (This property works as documented, with the exception of the prepare on EXECUTE STATEMENT, which can never be deferred for Cloudscape.)
- `resultSetHoldability` (Defaults to `CLOSE_CURSORS_AT_COMMIT` for Cloudscape. This is the only valid value for Network Server. Use the JDBC 3.0 API for setting holdability.)
- `securityMechanism` (Network Server supports these security mechanisms, the first being the default:  
`com.ibm.db2.jcc.DB2BaseDataSource.CLEAR_TEXT_PASSWORD_SECURITY`  
`com.ibm.db2.jcc.DB2BaseDataSource.ENCRYPTED_PASSWORD_SECURITY`



## Problems accessing a Sybase data source

What kind of problem are you having accessing your Sybase database?

- "Sybase Error 7713: Stored Procedure can only be executed in unchained transaction mode" error.
- "JZ0XS: The server does not support XA-style transactions. Please verify that the transaction feature is enabled and licensed on this server."
- A container managed persistence (CMP) enterprise bean is causing exceptions.

### "Sybase Error 7713: Stored Procedure can only be executed in unchained transaction mode" error

This error occurs when either:

- The JDBC attempts to put the connection in **autocommit(true)** mode.
- A stored procedure is not created in a compatible mode.

To fix the **autocommit(true)** mode problem, let the application change the connection to chained mode using the **Connection.setAutoCommit(false)** mode, or use a **set chained on** language command.

To resolve the stored procedure problem, use the `sp_procxmode procedure_name anymode` command.

### "JZ0XS: The server does not support XA-style transactions. Please verify that the transaction feature is enabled and licensed on this server."

This error occurs when XA-style transactions are attempted on a server that does not have Distributed Transaction Management (DTM) installed.

To resolve this problem, use the instructions in the Sybase Manual titled: *Using Adaptive Server Distributed Transaction Management Features* to enable Distributed Transaction Management (DTM). The main steps in this procedure are:

1. Install the DTM option.
2. Check the `license.dat` file to verify that the DTM option is installed.
3. Restart the license manager.
4. Enable DTM in ISQL.
5. Restart the ASE service.

### A container managed persistence (CMP) enterprise bean is causing exceptions

This error is caused by improper use of reserved words. Reserved words cannot be used as column names.

To correct this problem: Rename the variable to remove the reserved word. You can find a list of reserved words in the *Sybase Adaptive Server Enterprise Reference Manual; Volume 1: Building Blocks*, Chapter 4. This manual is available online at: <http://manuals.sybase.com/onlinebooks/group-as/asg1250e/refman>.

## Cannot access an enterprise bean from a servlet, a JSP file, a stand-alone program, or another client

What kind of error are you seeing?

- **javax.naming.NameNotFoundException: Name *name* not found in context "local" message** when access is attempted
- **BeanNotReentrantException** is thrown
- **CSITransactionRolledbackException / TransactionRolledbackException** is thrown

- Call fails, Stack trace beginning **EJSContainer E Bean method threw exception [exception\_name]** found in JVM log file.
- Call fails, **ObjectNotFoundException** or **ObjectNotFoundExceptionLocalException** when accessing stateful session EJB found in JVM log file.
- Attempt to start CMP EJB module fails with **javax.naming.NameNotFoundException: dataSourceName**
- 
- **Message BBOT0003W is issued**
- Symptom: **CNTR0001W: A Stateful SessionBean could not be passivated**

If the client is remote to the enterprise bean, which means, running in a different application server or as a stand-alone client, browse the logs of the application server hosting the enterprise bean as well as log files of the client.

### **ObjectNotFoundException or ObjectNotFoundExceptionLocalException when accessing stateful session EJB**

A possible cause of this problem is that the stateful session bean timed out and was removed by the container. This event must be addressed in the code, according to the EJB 2.0 specification (available at <http://java.sun.com/products/ejb/docs.html>), section 7.6.2, Dealing with exceptions.

### **Stack trace beginning "EJSContainer E Bean method threw exception [exception\_name]" found in JVM log file**

If the exception name indicates an exception thrown by an IBM class that begins with "com.ibm...", then search for the exception name within the information center, and in the online help as described below. If "exception name" indicates an exception thrown by your application, contact the application developer to determine the cause.

### **javax.naming.NameNotFoundException: Name name not found in context "local"**

A possible reason for this exception is that the enterprise bean is not local (not running in the same Java virtual machine [JVM] or application server) to the client JSP, servlet, Java application, or other enterprise bean, yet the call is to a "local" interface method of the enterprise bean. If access worked in a development environment but not when deployed to WebSphere Application Server, for example, it might be that the enterprise bean and its client were in the same JVM in development, but are in separate processes after deployment.

To resolve this problem, contact the developer of the enterprise bean and determine whether the client call is to a method in the local interface for the enterprise bean. If so, have the client code changed to call a remote interface method, or to promote the local method into the remote interface.

References to enterprise beans with local interfaces are bound in a name space local to the server process with the URL scheme of local:.

### **BeanNotReentrantException is thrown**

This problem can occur because client code (typically a servlet or JSP file) is attempting to call the same stateful SessionBean from two different client threads. This situation often results when an application stores the reference to the stateful session bean in a static variable, uses a global (static) JSP variable to refer to the

stateful SessionBean reference, or stores the stateful SessionBean reference in the HTTP session object. The application then has the client browser issue a new request to the servlet or JSP file before the previous request has completed.

To resolve this problem, ask the developer of the client code to review the code for these conditions.

### **CSITransactionRolledbackException / TransactionRolledbackException is thrown**

An enterprise bean container throws these high-level exceptions to indicate that an enterprise bean call could not successfully complete. When this exception is thrown, logs to determine the underlying cause.

Some possible causes include:

- The enterprise bean might throw an exception that was not declared as part of its method signature. The container is required to roll back the transaction in this case. Common causes of this situation are where the enterprise bean or code that it calls throws a `NullPointerException`, `ArrayIndexOutOfBoundsException`, or other Java runtime exception, or where a BMP bean encounters a JDBC error. The resolution is to investigate the enterprise bean code and resolve the underlying exception, or to add the exception to the problem method signature.
- A transaction might attempt to do additional work after being placed in a "Marked Rollback", "RollingBack", or "RolledBack" state. Transactions cannot continue to do work after they are set to one of these states. This situation occurs because the transaction has timed out which, often occurs because of a database deadlock. Work with the application database management tools or administrator to determine whether database transactions called by the enterprise bean are timing out.
- A transaction might fail on commit due to dangling work from local transactions. The local transaction encounters some "dangling work" during commit. When a local transactions encounters an "unresolved action" the default action is to "rollback". You can adjust this action to "commit" in the Application Assembly Tool (AAT). Open the enterprise bean .jar file (or the EAR file containing the enterprise bean) and select the Session Beans or Entity Beans object in the component tree on the left. The Unresolved Action property is on the IBM Extensions tab of the container properties.

### **Attempt to start EJB module fails with "javax.naming.NameNotFoundException dataSourceName\_CMP"exception**

This problem can occur because:

- When the DataSource resource was configured, container managed persistence was not selected.
  - To confirm this problem, in the administrative console, browse the properties of the data source given in the NameNotFoundException. On the Configuration panel, look for a check box labeled **Container Managed Persistence**.
  - To correct this problem, select the check box for **Container Managed Persistence**.
- If container managed persistence is selected, it is possible that the CMP DataSource could not be bound into the namespace.
  - Look for additional naming warnings or errors in the status bar, and in the hosting application server logs. Check any further naming-exception problems that you find by looking at the topic "Cannot look up an object hosted by WebSphere Application Server from a servlet, JSP file, or other client" on page 123.

## Message BBOT0003W is issued

Message BBOT0003W indicates a transaction timeout. The timeout might result in the abnormal termination of the servant where the transaction is executing.

- The default timeout value for enterprise bean transactions is 120 seconds. After this time, the transaction times out and the connection closes.
- If the transaction legitimately takes longer than the specified timeout period, on the administrative console:
  1. Go to **Manage Application Servers > server\_name**
  2. Select the **Transaction Service properties** page
  3. Increase the **Total transaction lifetime timeout** value
  4. **Save** the configuration

**Note:** z/OS will use the value you set for **Total transaction lifetime timeout** as the default transaction timeout setting. If you set a value for this property that is greater than the maximum transaction timeout value, z/OS will use the maximum transaction timeout value as the default.

## Symptom:CNTR0001W: A Stateful SessionBean could not be passivated

This error can occur when a Connection object used in the bean is not closed or nulled out.

To confirm this is the problem, look for an exception stack in the logs for the EJB container that hosts the enterprise bean, and looks similar to:

```
StatefulPassi W CNTR0001W:
A Stateful SessionBean could not be passivated: StatefulBean0
(BeanId(XXX#YYY.jar#ZZZ)),
state = PASSIVATING)
java.io.NotSerializableException: com.ibm.ws.rsadapter.jdbc.WSJdbcConnection
 at java.io.ObjectOutputStream.writeObject((Compiled Code))
 at java.io.ObjectOutputStream.writeObject(ObjectOutputStream.java(Compiled Code))
 at java.io.ObjectOutputStream.outputClassFields((Compiled Code))
 at java.io.ObjectOutputStream.defaultWriteObject((Compiled Code))
 at java.io.ObjectOutputStream.writeObject((Compiled Code))
 at java.io.ObjectOutputStream.writeObject(ObjectOutputStream.java(Compiled Code))
 at com.ibm.ejs.container.passivator.StatefulPassivator.passivate((Compiled Code))
 at com.ibm.ejs.container.StatefulBean0.passivate((Compiled Code))
 at com.ibm.ejs.container.activator.StatefulASActivationStrategy.atUnitOfWorkEnd
 ((Compiled Code))
 at com.ibm.ejs.container.activator.Activator.unitOfWorkEnd((Compiled Code))
 at com.ibm.ejs.container.ContainerAS.afterCompletion((Compiled Code))
```

where XXX,YYY,ZZZ is the Bean's name.

To correct this problem, the application must close all connections and set the reference to null for all connections. Typically this activity is done in the `ejbPassivate()` method of the bean. See the enterprise bean specification mandating this requirement, specifically section 7.4 in the EJB specification Version 2.0. Also, note that the bean must have code to reacquire these connections when the bean is reactivated. Otherwise, there are `NullPointerExceptions` when the application tries to reuse the connections.

## Cannot look up an object hosted by WebSphere Application Server from a servlet, JSP file, or other client

To resolve problems encountered when a servlet, JSP file, stand-alone application or other client attempts to access an enterprise bean, `ConnectionPool`, or other

named object hosted by WebSphere Application Server, you must first verify that the target server can be accessed from the client:

- From a command prompt on the client's server, enter "ping *server\_name*" and verify connectivity.
- Use the WebSphere Application Server administrative console to verify that the target resource's application server and, if applicable, EJB module or Web module, is started.

Continue only if there is no problem with connectivity and the target resource appears to be running.

What kind of error are you seeing?

- `NameNotFoundException` from JNDI lookup operation
- `CannotInstantiateObjectException` from JNDI lookup operation
- Message NMSV0610I appears in the server's log file, indicating that some Naming exception has occurred
- `OperationNotSupportedException` from JNDI Context operation.
- "WSVR0046E: Failed to bind" error, with Original exception: "org.omg.CosNaming.NamingContextPackage.AlreadyBound".
- `ConfigurationException` from "new InitialContext" operation or from a JNDI Context operation with a URL name.
- `ServiceUnavailableException` from "new InitialContext" operation.
- `CommunicationException` thrown from a "new InitialContext" operation.
- 

#### **NameNotFoundException from JNDI lookup operation**

If you encounter this exception in trying to access an enterprise bean, data source, messaging resource, or other resource:

- Browse the properties of the target object in the administrative console, and verify that the jndi name it specifies matches the JNDI name the client is using.
- If you are looking up an object that resides on a server different from the one from which the initial context was obtained, you must use the fully qualified name.
  - If access is from another server object such as a servlet accessing an enterprise bean and you are using the default context, not specifying the fully qualified JNDI name, you may get this error if the object is being hosted on a different server.
  - If access is from a stand-alone client, it may be that the object you are attempting access is on a server different from the server from which you obtained the initial context.

To correct this problem, use the fully-qualified JNDIname:

- If the object is in a single server:  
`cell/nodes/nodeName/servers/serverName/jndiName`. Objects are not supported in this release.
- If the object is on a server cluster: `cell/clusters/clusterName/jndiName`.

#### **CannotInstantiateObjectException from JNDI lookup operation**

If you encounter this exception in trying to access an enterprise bean, data source, messaging resource, or other resource, possible causes include:

- A serialized Java object is being looked up, but the necessary classes required to deserialize it are not in the runtime environment.
- A Reference object is being looked up, and the associated factory used to process it as part of the lookup processing is failing.

To determine the precise cause of the problem:

- Look in the logs of the server hosting the target resource. Look for exceptions immediately preceding the `CannotInstantiateObjectException`. If it is a `java.lang.NoClassDefFoundError` or `java.lang.ClassNotFoundException`, make sure the class referenced in the error message can be located by the class loader.
- Print out the stack trace for the root cause and look for the factory class. It will be called by `javax.naming.NamingManager.getObjectInstance()`. The reason for the failure will depend on the factory implementation, and may require you to contact the developer of the factory class.

### **Message NMSV0610I appears in the server's log file, indicating that some Naming exception has occurred**

This error is informational only and is provided in case the exception is related to an actual problem. Most of the time, it is not. If it is, the log file should contain adjacent entries to provide context.

- If no problems are being experienced, ignore this message. Also ignore the message if the problem you are experiencing does not appear to be related to the exception being reported and if there are no other adjacent error messages in the log.
- If a problem is being experienced, look in the log for underlying error messages.
- The information provided in message NMSV0610I can provide valuable debug data for other adjacent error messages posted in response to the Naming exception that occurred.

### **OperationNotSupportedException from JNDI Context operation**

This error has two possible causes:

- An update operation, such as a bind, is being performed with a name that starts with `"java:comp/env"`. This context and its subcontexts are read-only contexts.
- A Context bind or rebind operation of a non-CORBA object is being performed on a remote name space that does not belong to WebSphere Application Server. Only CORBA objects can be bound to these `CosNaming` name spaces.

To determine which of these errors is causing the problem, check the full exception message.

### **WSVR0046E: Failed to bind, ejb/jndiName: ejb/jndiName. Original exception : org.omg.CosNaming.NamingContextPackage.AlreadyBound**

This error occurs two enterprise bean server applications were installed on the same server such that a binding name conflict occurred. That is, a `jndiName` value is the same in the two applications' deployment descriptors. The error will surface during server startup when the second application using that `jndiName` value is started.

To verify that this is the problem, examine the deployment descriptors for all enterprise bean server applications running in the server in search for a `jndiName` that is specified in more than one enterprise bean application.

To correct the problem, change any duplicate `jndiName` values to ensure that each enterprise bean in the server process is bound with a different name.



### **ConfigurationException from "new InitialContext" operation or from a JNDI Context operation with a URL name**

If you are attempting to obtain an initial JNDI context, a configuration exception can occur because an invalid JNDI property value was passed to the InitialContext constructor. This includes JNDI properties set in the System properties or in some jndi.properties file visible to the class loader in effect. A malformed provider URL is the most likely property to be incorrect. If the JNDI client is being run as a thin client such that the CLASSPATH is set to include all of the individual jar files required, make sure the .jar file containing the properties file `com/ibm/websphere/naming/jndiprovider.properties` is in the CLASSPATH.

If the exception is occurring from a JNDI Context call with a name in the form of a URL, the current JNDI configuration may not be set up properly so that the required factory class name cannot be determined, or the factory may not be visible to the class loader currently in effect. If the name is a Java: URL, the JNDI client must be running in a J2EE client or server environment. That is, the client must be running in a container.

Check the exception message to verify the cause.

If the exception is being thrown from the InitialContext constructor, correct the property setting or the CLASSPATH.

If the exception is being thrown from a JNDI Context method, make sure the property `java.naming.factory.url.pkgs` includes the package name for the factory required for the URL scheme in the name. URL names with the Java scheme can only be used while running in a container.

### **ServiceUnavailableException from "new InitialContext" operation**

This exception indicates that some unexpected problem occurred while attempting to contact the name server to obtain an initial context. The ServiceUnavailableException, like all NamingException objects, can be queried for a root cause. Check the root cause for more information. It is possible that some of the problems described for CommunicationExceptions may also result in a ServiceUnavailableException.

Since this exception is triggered by an unexpected error, there is no probable cause to confirm. If the root cause exception does not indicate what the probable cause is, investigate the possible causes listed for CommunicationExceptions.

### **CommunicationException thrown from a "new InitialContext" operation**

The name server identified by the provider URL cannot be contacted to obtain the initial JNDI context. There are many possible causes for this problem, including:

- The host name or port in the provider URL is incorrect.
- The host name cannot be resolved into an IP address by the domain name server, or the IP address does not match the IP address which the server is actually running under.
- A firewall on the client or server is preventing the port specified in the provider URL from being used.

To correct this problem:

- Make sure the provider URL and the network configurations on the client and server machines are correct.

- Make sure the host name can be resolved into an IP address which can be reached by the client machine. You can do this using the ping command.
- If you are running a firewall, make sure that use of the port specified in the provider URL will be allowed.

## Access problems after enabling security

What kind of error are you seeing?

- I cannot access all or part of the administrative console or use the wsadmin tool after enabling security
- I cannot access a Web page after enabling security
- Authentication error accessing a Web page
- Authorization error accessing a Web page
- The client cannot access an enterprise bean after enabling security
- The client never gets prompted when accessing a secured enterprise bean
- I cannot stop an application server, node manager, or node after enabling security
- After enabling single signon, I cannot log on to the administrative console.

### I cannot access all or part of the administrative console or use the wsadmin tool after enabling security

- If you cannot access the administrative console, or view and update certain objects, look in the logs of the application server which hosts the administrative console page for a related error message.
- You might not have authorized your ID for administrative tasks. This problem is indicated by errors such as:
  - [8/2/02 10:36:49:722 CDT] 4365c0d9 RoleBasedAuth A SECJ0305A: Role based authorization check failed for security name MyServer/myUserId, accessId MyServer/S-1-5-21-882015564-4266526380-2569651501-1005 while invoking method getProcessType on resource Server and module Server.
  - Exception message: "ADMN0022E: Access denied for the getProcessType operation on Server MBean"
  - When running the command: `wsadmin -username j2ee -password j2ee: WASX7246E: Cannot establish "SOAP" connection to host "BIRKT20" because of an authentication failure. Please ensure that user and password are correct on the command line or in a properties file.`

To grant an ID administrative authority, from the administrative console, click **System Administration > Console Users** and validate that the ID is a member. If it is not, add the ID with at least monitor access privileges, for read-only access.

- Check that the *enable\_trusted\_application* flag is set to true. To check the *enable\_trusted\_application* flag, from the Administrative Console, click **Security > Global Security > Custom Properties > Enable Trusted Application** and verify that it is set to true.

### I cannot access a Web page after enabling security

When secured resources are not accessible, probable causes include:

- Authentication errors - WebSphere Application Server security cannot identify the ID of the person or process. Symptoms of authentication errors include:

On a Netscape browser:

- Authorization failed. Retry? message displays after an attempt to log in.
- Accepts any number of attempts to retry login and displays Error 401 message when Cancel is clicked to stop retry.
- A typical browser message displays: Error 401: Basic realm='Default Realm'.



On an Internet Explorer browser:

- Login prompt displays again after an attempt to log in.
- Allows three attempts to retry login.
- Displays Error 401 message after three unsuccessful retries.
- Authorization errors - the security function has identified the requesting person or process as not authorized to access the secured resource. Symptoms of authorization errors include:
  - Netscape browser: "Error 403: AuthorizationFailed" message is displayed.
  - Internet Explorer:
    - "You are not authorized to view this page" message is displayed.
    - "HTTP 403 Forbidden" error is also displayed.
- SSL errors - WebSphere Application Server security uses Secure Socket Layer (SSL) technology internally to secure and encrypt its own communication, and incorrect configuration of the internal SSL settings can cause problems. Also you might have enabled SSL encryption for your own Web application or enterprise bean client traffic which, if configured incorrectly, can cause problems regardless of whether WebSphere Application Server security is enabled.
  - SSL related problems are often indicated by error messages which contain a statement such as: ERROR: Could not get the initial context or unable to look up the starting context.Exiting. followed by  
javax.net.ssl.SSLHandshakeException
  - System SSL failures are indicated by error messages like the following:

```
Trace: 2003/08/11 19:53:40.682 01 t=9D49C0 c=UNK key=S2 (0E012048)
Description: Failure Exit from -> SecurityManager::secureSocketInit
Socket descriptor: 211
gsk_secure_socket_init: 410
```

### The client cannot access an enterprise bean after enabling security

If client access to an enterprise bean fails after security is enabled:

- Review the steps for securing and granting access to resources.
- Browse the server logs for errors relating to enterprise bean access and security. Look up any errors in the message table.

Errors similar to *Authorization failed for /UNAUTHENTICATED while invoking resource securityName:/UNAUTHENTICATED;accessId:UNAUTHENTICATED not granted any of the required roles* indicate that:

- An unprotected servlet or JavaServer Page (JSP) file accessed a protected enterprise bean. When an unprotected servlet is accessed, the user is not prompted to log in and the servlet runs as UNAUTHENTICATED. When the servlet makes a call to an enterprise bean that is protected the servlet fails.

To resolve this problem, secure the servlet that is accessing the protected enterprise bean. Make sure the *runAs* property for the servlet is set to an ID that can access the enterprise bean.

- An unauthenticated Java client program is accessing an enterprise bean resource that is protected. This situation can happen if the file read by the *sas.client.props* properties file used by the client program does not have the *securityEnabled* flag set to **true**.

To resolve this problem, make sure that the *sas.client.props* file on the client side has its *securityEnabled* flag set to **true**.

Errors similar to **Authorization failed for *valid\_user* while invoking resource securityName:/username;accessId:xxxxxx not granted any of the required roles** indicate that a client attempted to access a secured enterprise bean resource, and the supplied user ID is not assigned the required roles for that enterprise bean.

- Check that the required roles for the enterprise bean resource are accessed. View the required roles for the enterprise bean resource in the deployment descriptor of the Web resource.
- Check the authorization table and make sure that the user or the group that the user belongs to is assigned one of the required roles. You can view the authorization table for the application that contains the enterprise bean resource using the administrative console.
- If you are using LOCALOS and SAFAuthorization, check the SAF EJBROLES setup. Specifically, verify that
  - the EJBROLE class has been activated.
  - The role has been defined to SAF.
  - The userid has been permitted to the role.
  - The class was refreshed after the permit operation.

### **Client program never gets prompted when accessing secured enterprise bean**

Even though it appears security is enabled and an enterprise bean is secured, it can happen that the client executes the remote method without prompting. If the remote method is protected, an authorization failure results. Otherwise, execute the method as an unauthenticated user.

Possible reasons for this problem include:

- The server with which you are communicating might not have security enabled. Check with the WebSphere Application Server administrator to ensure that the server security is enabled. Access the global security settings from within the Security section of the administrative console.
- The client does not have security enabled in the sas.client.props file. Edit the sas.client.props file to ensure the property **com.ibm.CORBA.securityEnabled** is set to true.
- The client does not have a ConfigURL specified. Verify that the property **com.ibm.CORBA.ConfigURL** is specified on the command line of the Java client, using the -D parameter.
- The specified ConfigURL has an invalid URL syntax, or the sas.client.props that is pointed to cannot be found. Verify that the **com.ibm.CORBA.ConfigURL** property is valid, for example, similar to the file:/WebSphere/AppServer/properties/sas.client.props file on Windows systems. Check the Java documentation for a description of URL formatting rules. Also, validate that the file exists at the specified path.

### **Cannot stop an application server, node manager, or node after enabling security**

If you use command line utilities to stop WebSphere Application Server processes, apply additional parameters after enabling security to provide authentication and authorization information.

### **After enabling single signon, I cannot log on to the administrative console**

This problem occurs when single signon (SSO) is enabled, and you attempt to access the administrative console using the short name of the server, for example `http://myserver:9090/admin`. The server accepts your user ID and password, but returns you to the log on page instead of the administrative console.

To correct this problem, use the fully qualified host name of the server, for example `http://myserver.mynetwork.mycompany.com:9090/admin`.

## Errors after enabling security

What kind of error are you seeing?

- Authentication error accessing a Web page
- Authorization error accessing a Web page
- Error Message: SECJ0314E: Current Java 2 security policy reported a potential violation
- MSGS0508E: The JMS Server security service was unable to authenticate user ID; error displayed in SystemOut.log when starting an application server
- Error Message: SECJ0237E: One or more vital LTPAServerObject configuration attributes are null or not available after enabling security and starting the application server.
- Error Message: SECJ0336E: Authentication failed for user {0} because of the following exception {1}

### Authentication error accessing a Web page

Possible causes for authentication errors include:

- **Incorrect user name or passwords.** Check the user name and password and make sure they are correct.
- **Security configuration error : User registry type is not set correctly.** Check the user registry property in global security settings in the administrative console. Verify that it is the intended user registry.
- **Internal program error.** If the client application is a Java standalone program, this program might not gather or send credential information correctly.

### Authorization error accessing a Web page

If a user who should have access to a resource does not, there is probably a missing configuration step. Review the steps for securing and granting access to resources.

Specifically:

- Check required roles for the accessed Web resource.
- Check the authorization table to make sure that the user, or the groups to which the user belongs, is assigned to one of the required roles.
- View required roles for the Web resource in the deployment descriptor of the Web resource.
- View the authorization table for the application that contains the Web resource, using the administrative console.
- Test with a user who is granted the required roles, to see if the user can access the problem resources.
- If the problem user is required to have one or more of the required roles, use the administrative console to assign that user to required roles. Then stop and restart the application.

### Error Message: SECJ0314E: Current Java 2 security policy reported a potential violation on server

If you find errors on your server similar to:

```
Error Message: SECJ0314E: Current Java 2 Security policy reported
a potential violation of Java 2 Security Permission.
Please refer to Problem Determination Guide for further information.
{0}Permission\:{1}Code\:{2}{3}Stack Trace\:{4}Code Base Location\:{5}
```

The Java security manager `checkPermission()` method has reported an exception, `SecurityException`.

**The reported exception might be critical to the secure system.** Turn on security trace to determine the potential code that might have violated the security policy. Once the violating code is determined, verify if the attempted operation is permitted with respect to Java 2 Security, by examining all applicable Java 2 security policy files and the application code.

A more detailed report is enabled by either configuring RAS trace into debug mode, or specifying a Java property.

- Specify the following property in the **Application Servers > server name > ProcessDefinition > Java Virtual Machine** panel from the administrative console in the **Generic JVM arguments** panel:
  - Add the run-time flag `java.security.debug`
  - Valid values:
    - access** Print all debug information including: required permission, code, stack, and code base location.
    - stack** Print debug information including: required permission, code, and stack.
    - failure** Print debug information including: required permission and code.

For a review of Java security policies and what they mean, see the Java 2 Security documentation at <http://java.sun.com/j2se/1.3/docs/guide/security/index.html>.

**Tip:** If the application is running with a Java Mail API, this message might be benign. You can update the *installed Enterprise Application root/META-INF/was.policy* file to grant the following permissions to the application:

- `permission java.io.FilePermission "${user.home}${/}.mailcap", "read";`
- `permission java.io.FilePermission "${user.home}${/}.mime.types", "read";`
- `permission java.io.FilePermission "${java.home}${/}lib${/}mailcap", "read";`
- `permission java.io.FilePermission "${java.home}${/}lib${/}mime.types", "read";`

**Error message: MSGS0508E: The JMS Server security service was unable to authenticate user ID:" error displayed in SystemOut.log when starting an application server**

This error can result from installing the JMS messaging API sample and then enabling security. You can follow the instructions in the Configure and Run page of the corresponding JMS sample documentation to configure the sample to work with WebSphere Application Server security.

You can verify the installation of the message-driven bean sample by launching the installation program, selecting **Custom**, and browsing the components which are already installed in the **Select the features you like to install** panel. The JMS sample is shown as **Message-Driven Bean Sample**, under **Embedded Messaging**.

You can also verify this installation by using the administrative console to open the properties of the application server which contains the samples. Select **MDBSamples** and click **uninstall**.

**Error message: SECJ0237E: One or more vital LTPAServerObject configuration attributes are null or not available after enabling security and starting the application server.**

This error message can result from selecting LTPA as the authentication mechanism, but not generating the LTPA keys. The LTPA keys encrypt the LTPA token.

To resolve this problem:

1. Click **System Administration > Console users > LTPA**
2. Enter a password, which can be anything.
3. Enter the same password in **Confirm Password**.
4. Click **Apply**.
5. Click **Generate Keys**.
6. Click on **Save**.

**The exception `AccessControlException`, is reported in the `SystemOut.log`**

The problem is related to the Java 2 Security feature of WebSphere Application Server, the API-level security framework that is implemented in WebSphere Application Server Version 5. An exception similar to the following example displays. The error message and number can vary.

```
E SRVE0020E: [Servlet Error]-[validator]: Failed to load servlet:
java.security.AccessControlException:
access denied
(java.io.FilePermission
\WebSphere\AppServer\installedApps\maeda\adminconsole.ear\adminconsole.war\
WEB-INF\validation.xml read)
```

For an explanation of Java 2 security, how and why to enable or disable it, how it relates to policy files, and how to edit policy files, see the Java 2 security topic in the information center navigation. The topic explains that Java 2 security is not only used by this product, but developers can also implement it for their business applications. Administrators might need to involve developers, if this exception is thrown when a client tries to access a resource hosted by WebSphere Application Server.

Possible causes of these errors include:

- Syntax errors in a policy file.
- Syntax errors in permission specifications in the `ra.xml` file bundled in a `.rar` file. This case applies to resource adapters that support connector access to CICS or other resources.
- An application is missing the specified permission in a policy file, or in permission specifications in an `ra.xml` file bundled in a `.rar` file
- The class path is not set correctly, preventing the permissions for the `resource.xml` file for SPI from being correctly created.
- A library called by an application, or the application, is missing a `doPrivileged` block to support access to a resource.
- Permission is specified in the wrong policy file.

To resolve these problems:

- Check all of the related policy files to verify that the permission shown in the exception, for example `java.io.FilePermission`, is specified.
- Look for a related `ParserException` in the `SystemOut.log` file which reports the details of the syntax error. For example: `SECJ0189E: Caught ParserException while creating template for application policy`

\WAS\config\cells\xxx\nodes\xxx\app.policy. The exception is com.ibm.ws.security.util.ParserException: line 18: expected ';', found 'grant'

- Look for a message similar to: SECJ0325W: The permission *permission* specified in the policy file is unresolved.
  - Check the call stack to determine which method does not have the permission. Identify the class path of this method. If it is hard to identify the method, enable the Java2 security Report.
- The report shows:

**Permission**

the missing permission.

**Code** which method has the problem.

**Stack Trace**

where the access violation occurred.

**CodeBaseLocation**

the detail of each stack frame.

Usually, Permission and Code are enough to identify the problem. The following example illustrates a report:

Permission:

```
\WebSphere\AppServer\logs\server1\SystemOut_02.08.20_11.19.53.log :
access denied (java.io.FilePermission
```

```
\WebSphere\AppServer\logs\server1\SystemOut_02.08.20_11.19.53.log delete)
```

Code:

```
com.ibm.ejs.ras.RasTestHelper$7 in
{
file:/WebSphere/AppServer/installedApps/maeda/JrasFVTApp.ear/RasLib.jar}
```

Stack Trace:

```
java.security.AccessControlException: access denied (java.io.FilePermission
```

```
\WebSphere\AppServer\logs\server1\SystemOut_02.08.20_11.19.53.log delete)
at java.security.AccessControlContext.checkPermission
 (AccessControlContext.java(Compiled Code))
at java.security.AccessController.checkPermission
 (AccessController.java(Compiled Code))
at java.lang.SecurityManager.checkPermission
 (SecurityManager.java(Compiled Code))
 .
```

Code Base Location:

```
com.ibm.ws.security.core.SecurityManager :
```

```
file:/WebSphere/AppServer/lib/securityimpl.jar
ClassLoader: com.ibm.ws.bootstrap.ExtClassLoader
Permissions granted to CodeSource
(
file:/WebSphere/AppServer/lib/securityimpl.jar <no certificates>
{
 (java.util.PropertyPermission java.vendor read);
 (java.util.PropertyPermission java.specification.version read);
 (java.util.PropertyPermission line.separator read);
 (java.util.PropertyPermission java.class.version read);
 (java.util.PropertyPermission java.specification.name read);
 (java.util.PropertyPermission java.vendor.url read);
 (java.util.PropertyPermission java.vm.version read);
 (java.util.PropertyPermission os.name read);
 (java.util.PropertyPermission os.arch read);
}
(This list continues.)
```



- If the method is SPI, check the resources.xml file to ensure that the class path is correct.
- To confirm that all of the policy files are loaded correctly, or what permission each class path is granted, enable the trace with **com.ibm.ws.security.policy.\*=all=enabled**. All loaded permissions are listed in the trace.log file. Search for the app.policy, was.policy and ra.xml files. To check the permission list for a class path, search for **Effective Policy for classpath**.
- If there are any syntax errors in the policy file or ra.xml file, correct them with the policytool. Avoid editing the policy manually, because syntax errors can result.
- If a permission is listed as Unresolved, it does not take effect. Verify that the specified permission name is correct.
- If the class path specified in the resource.xml file is not correct, correct it.
- If a required permission does not exist in either the policy files or the ra.xml file, examine the application code to see if you need to add this permission. If so, add it to the proper policy file or ra.xml file.
- If the permission should not be granted outside of the specific method that is accessing this resource, modify the code needs to use a doPrivileged block.
- If this permission does exist in a policy file or a ra.xml file and they were loaded correctly, but the class path still does not have the permission in its list, the location of the permission might not be correct. Read Java 2 Security in the information center navigation carefully to determine in which policy file or ra.xml file that permission should be specified.

**Tip:** If the application is running with the Java Mail API, you can update the *installed Enterprise Application root/META-INF/was.policy* file to grant the following permissions to the application:

- permission java.io.FilePermission "\${user.home}\${/}.mailcap", "read";
- permission java.io.FilePermission "\${user.home}\${/}.mime.types", "read";
- permission java.io.FilePermission "\${java.home}\${/}lib\${/}mailcap", "read";
- permission java.io.FilePermission "\${java.home}\${/}lib\${/}mime.types", "read";

**Error Message: SECJ0336E: Authentication failed for user {0} because of the following exception {1}**

This error message results if the user ID indicated is not found in the LDAP user registry. To resolve this problem:

1. Verify that your user ID and password are correct.
2. Verify that the user ID exists in the registry.
3. Verify that the base distinguished name (DN) is correct.
4. Verify that the user filter is correct.
5. Verify that the bind DN and the password for the bind DN are correct. If the bind DN and password are not specified, add the missing information and retry.
6. Verify that the host name and LDAP type are correct.

Consult with the administrator of the user registry if the problem persists.

## Errors after configuring or enabling Secure Sockets Layer

This article explains various problems you might encounter after configuring or enabling Secure Sockets Layer (SSL).



## Stopping the deployment manager after configuring Secure Sockets Layer **5.1+**

After configuring the Secure Sockets Layer repertoires, if you stop the deployment manager without also stopping the node agents, you might receive the following error message when you restart the deployment manager:

```
ADMU0509I: The server "nodeagent" cannot be reached. It appears to be stopped.
ADMU0211I: Error details may be seen in the file:
 /opt/WebSphere/AppServer/logs/nodeagent/stopServer.log
```

The error occurs because the deployment manager did not propagate the new SSL certificate to the node agents. Thus, the node agents are using an older certificate files than the deployment manager and the certificate files are incompatible. To work around this problem, you must manually stop the node agent and deployment manager processes. To kill the processes on the z/OS platform, use the MVS console and type `c process_name`.

### Accessing resources using HTTPS

If you are unable to access resources using a Secure Sockets Layer Web address (beginning with `https:`), or encounter error messages which indicate SSL problems, verify that your HTTP server is configured correctly for SSL by browsing the welcome page of the HTTP server using SSL by entering the URL:  
**`https://hostname`**.

If the page works with HTTP, but not HTTPS, the problem is with the HTTP server.

Refer to the documentation for your HTTP server for instructions on correctly enabling SSL. If you are using the IBM HTTP Server or Apache, go to: <http://www.ibm.com/software/webservers/httpservers/library.html>. Click **Frequently Asked Questions > SSL**.

If the HTTP server handles SSL-encrypted requests successfully, or is not involved (for example, traffic flows from a Java client application directly to an enterprise bean hosted by the WebSphere Application Server, or the problem appears only after enabling WebSphere Application Server security), what kind of error are you seeing?

### System SSL

See *z/OS System Secure Sockets Layer Programming SC24-5901* for information on using the System Secure Sockets Layer (SSL) callable services programming interfaces.

- `javax.net.ssl.SSLHandshakeException` - The client and server could not negotiate the desired level of security. Reason: handshake failure
- `javax.net.ssl.SSLHandshakeException` - The client and server could not negotiate the desired level of security. Reason: unknown certificate
- `javax.net.ssl.SSLHandshakeException` - The client and server could not negotiate the desired level of security. Reason: bad certificate
- error when programmatically creating a credential.

## **javax.net.ssl.SSLHandshakeException - The client and server could not negotiate the desired level of security. Reason: handshake failure**

If you see a Java exception stack similar to the following example:

```
[Root exception is org.omg.CORBA.TRANSIENT:
CAUGHT_EXCEPTION_WHILE_CONFIGURING_SSL_CLIENT_SOCKET:
JSSL0080E: javax.net.ssl.SSLHandshakeException -
The client and server could not negotiate the desired level of
security. Reason: handshake failure:host=MYSERVER,port=1079
minor code: 4942F303 completed: No] at
com.ibm.CORBA.transport.TransportConnectionBase.connect
(TransportConnectionBase.java:NNN)
```

Some possible causes are:

- Not having common ciphers between the client and server.
- Not specifying the correct protocol.

To correct these problems:

1. Review the SSL settings. Click **WebSphere Administrative Console Security Settings > SSL Configuration Repertoires > DefaultSSLSettings** (or other named SSL settings).
2. Select the **Secure Sockets Layer (SSL)** option from the Additional Properties menu. You can also browse the file manually by viewing the *install\_dir/properties/sas.client.props* file.
3. Check the property specified by the *com.ibm.ssl.protocol* file to determine which protocol is specified.
4. Check the cipher types specified by the *com.ibm.ssl.enabledCipherSuites*. You might want to add more cipher types to the list. To see which cipher suites are currently enabled: Go to the properties page of the SSL settings as described above, and look for the **Cipher Suites** property. To see the list of all possible cipher suites, go to the properties page of the SSL settings as described above, then view the online help for that page. From the help page, click **Configure additional SSL settings**.
5. Correct the protocol or cipher problem by using a different client or server protocol and cipher selection. Typical protocols are SSL or SSLv3.

## **javax.net.ssl.SSLHandshakeException: unknown certificate**

If you see a Java exception stack similar to the following example, it might be caused by not having the personal certificate for the server in the client truststore file:

```
ERROR: Could not get the initial context or unable
to look up the starting context. Exiting. Exception received:
javax.naming.ServiceUnavailableException: A communication failure
occurred while attempting to obtain an initial context using the
provider url: "corbaloc:iiop:localhost:2809". Make sure that the
host and port information is correct and that the server identified
by the provider url is a running name server. If no port number is
specified, the default port number 2809 is used. Other possible
causes include the network environment or workstation network
configuration. [Root exception is org.omg.CORBA.TRANSIENT:
CAUGHT_EXCEPTION_WHILE_CONFIGURING_SSL_CLIENT_SOCKET: JSSL0080E:
javax.net.ssl.SSLHandshakeException - The client and server could
not negotiate the desired level of security. Reason: unknown
certificate:host=MYSERVER,port=1940 minor code: 4942F303 completed: No]
```

To correct this problem:

1. Check the client truststore file to determine if the signer certificate from the server personal certificate is there. For a self-signed server personal certificate,

the signer certificate is the public key of the personal certificate. For a certificate authority signed server personal certificate, the signer certificate is the root CA certificate of the CA that signed the personal certificate.

2. Add the server signer certificate to the client truststore file.

### **javax.net.ssl.SSLHandshakeException: bad certificate**

If you see a Java exception stack similar to the following example, it can be caused by having a personal certificate in the client keystore used for SSL mutual authentication but not having extracted the signer certificate into the server truststore file so that the server can trust it whenever the SSL handshake is made:

```
ERROR: Could not get the initial context or unable
to look up the starting context. Exiting. Exception received:
javax.naming.ServiceUnavailableException: A communication
failure occurred while attempting to obtain an initial context
using the provider url: "corbaloc:iiop:localhost:2809". Make
sure that the host and port information is correct and that
the server identified by the provider url is a running name
server. If no port number is specified, the default port number
2809 is used.Other possible causes include the network
environment or workstation network configuration. [Root
exception is org.omg.CORBA.TRANSIENT: CAUGHT_EXCEPTION_
WHILE_CONFIGURING_SSL_CLIENT_SOCKET: JSSL0080E:
javax.net.ssl.SSLHandshakeException - The client and server
could not negotiate the desired level of security. Reason:
bad certificate: host=MYSERVER,port=1940 minor code: 4942F303
completed: No]
```

To verify this problem, check the server truststore file to determine if the signer certificate from the client personal certificate is there. For a self-signed client personal certificate, the signer certificate is the public key of the personal certificate. For a certificate authority signed client personal certificate, the signer certificate is the root CA certificate of the CA that signed the personal certificate.

To correct this problem, add the client signer certificate to the server truststore file.

## **Errors in messaging**

What kind of problem are you seeing?

- javax.jms.JMSEException: MQJMS2008: failed to open MQ queue in JVM log.

### **javax.jms.JMSEException: MQJMS2008: failed to open MQ queue in JVM log**

This error can occur when the MQ queue name is not defined in the internal Java Messaging Service (JMS) server queue names list. This problem can occur if a WebSphere Application Server queue destination is created, without adding the queue name to the internal JMS server queue names list.

To resolve this problem:

1. Open the WebSphere Application Server administrative console.
2. Click **Servers > Manage Application Servers > server\_name > Server Components > JMS Servers**.
3. Add the queue name to the list.
4. Save the changes and restart the server.

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

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

## Errors returned to a client sending a SOAP request

What kind of problem are you seeing?

- `SOAPException: faultCode=SOAP-ENV:Client; msg=Error opening socket; java.net.ConnectException: Connection refused: connect`
- `javax.security.cert.CertPathBuilderException: No end-entity certificate matching the selection criteria could be found.`

If none of these errors match the one you see:

- Browse the application server logs.
- Look up any error or warning messages in the message table.

**`SOAPException: faultCode=SOAP-ENV:Client; msg=Error opening socket; java.net.ConnectException: Connection refused: connect`**

The most likely cause of this refused connection is that it was sent to the default port, 80, and an HTTP server is not installed or configured.

To verify this situation, send the message directly to the SOAP port; for example, to `http://hostname:9080`. If the message is sent correctly, there are two ways to resolve the problem:

- Continue specifying port 9080 on SOAP requests.
- If an HTTP server is not installed, install one and the associated plug-in component.
- If an HTTP server is installed:
  - Regenerate the HTTP plug-in configuration in the administrative console by clicking **Environment > Update WebServer Plugin**, and restarting the HTTP server.
  - If the problem persists, view the HTTP server access and error logs, as well as the `install_dir/logs/http_plugin.log` file for more information.

**`javax.security.cert.CertPathBuilderException: No end-entity certificate matching the selection criteria could be found`**

This error usually indicates that new or updated security keys are needed. The security key files are:

- `SOAPclient`
- `SOAPserver`
- `sslserver.p12`

In an installed application, these files are located in the:

`install_dir/installedApps/application_name.ear/soapsec.war/key/` directory. After replacing these files, you must stop and restart the application.

To replace these files in a SOAP-enabled application that is not yet installed:

- Expand the `application_name.ear` file.
- Expand the `soapsec.war` file.
- Replace the security key files in the `key/` directory.
- After you replace these files, install the application and restart the server.

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

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

## Errors connecting to WebSphere MQ and creating WebSphere MQ queue connection factory

The following exception may occur when trying to create the MDBListener instance:

```
6/23/03 22:45:58:232 CDT] 673106a8 MsgListenerPo W WMSG0049E:
Failed to start MDB PSSampleMDB against listener port SamplePubSubListenerPort
[6/23/03 22:47:58:289 CDT] 673106a8 FreePool E J2CA0046E:
Method createManagedConnctionWithMCWrapper caught an exception
during creation of the ManagedConnection for resource
JMS$SampleJMSQueueConnectionFactory, throwing ResourceAllocationException.
Original exception: javax.resource.spi.ResourceAdapterInternalException:
 createQueueConnection failed
com.ibm.mq.MQException: MQJE001: An MQException occurred:
Completion Code 2, Reason 2009
MQJE003: IO error transmitting message buffer at
com.ibm.mq.MQManagedConnectionJ11. (MQManagedConnectionJ11.java:239)
```

This problem occurs because the MQ manager userid does not have write access to the /tmp directory. To correct this problem, before you use a Jacl procedure to configure WebSphere Application Server resources and install an application:

1. Ensure that all applications have write access to /tmp directory. Use the `chmod 1777` command on the directory if necessary.
2. Create another subdirectory under /tmp (for example, /tmp/mydir). Use this directory as a "working directory" for the Jacl.
3. Restart the server.

Applications that use messaging on startup should start successfully.

---

## Troubleshooting application runtime and management problems

Select the problem you are having with running or managing deployed code for WebSphere Application Server:

- I have problems bringing up or using the administrative console.
- I have problems starting or using the `wsadmin` command prompt.
- My Web module or application server dies or hangs.
- I get errors trying to configure and enable security.
- I cannot seem to distribute the workload across clustered servers.
- There are problems setting up the multiserver Deployment Manager environment.
- I cannot uninstall or remove a node or application server.
- I have problems creating or using HTTP sessions.
- I have problems using tracing, logging, logfiles, or other troubleshooting features.
- I get errors connecting to the administrative console from a Netscape browser.

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

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

## Problems starting or using the wsadmin command

What kind of problem are you having?

- "WASX7023E: Error creating "SOAP" connection to host" or similar error trying to launch wsadmin command line utility.
- "com.ibm.bsf.BSFException: error while eval'ing Jacl expression: no such method "<command name>" in class com.ibm.ws.scripting.AdminConfigClient" returned from wsadmin command.
- WASX7022E returned from running "wsadmin -c ..." command, indicating invalid command.
- com.ibm.ws.scripting.ScriptingException: WASX7025E: String "" is malformed; cannot create ObjectName.

If you do not see your problem here:

- If you are not able to enter wsadmin command mode, try running **wsadmin -c "\$Help wsadmin"** for help in verifying that you are entering the command correctly.
- If you can get the wsadmin command prompt, enter **\$Help help** to verify that you are using specific commands correctly.
- wsadmin commands are a superset of Jacl (Java Command Language), which is in turn a Java-based implementation of the Tcl command language. For details on Jacl syntax beyond wsadmin commands, refer to the Tcl developers' site, <http://www.tcl.tk>. For specific details relating to the Java implementation of Tcl, refer to <http://www.tcl.tk/software/java>.
- Browse the *install\_dir/logs/wsadmin.traceout* file for clues.
  - Keep in mind that wsadmin.traceout is refreshed (existing log records are deleted) whenever a new wsadmin session is started.
  - If the error returned by wsadmin does not seem to apply to the command you entered, for example, you receive WASX7023E, stating that a connection could not be created to host "myhost," but you did not specify "-host myhost" on the command line, examine the properties files used by wsadmin to determine what properties are specified. If you do not know what properties files were loaded, look for the WASX7326I messages in the wsadmin.traceout file; there will be one of these messages for each properties file loaded.

### "WASX7023E: Error creating "SOAP" connection to host" or similar error trying to launch wsadmin command line utility

By default, the wsadmin utility attempts to connect to an application server at startup. This is because some commands act upon running application servers. This error indicates that no connection could be established.

To resolve this problem:

- If you are not sure whether an application server is running, start it by entering **startserver.sh server short name** from the command prompt. If the server is already running, you will see an error similar to "ADMU3027E: An instance of the server is already running".
- If you are running a z/OS configuration, you will first need to start the deployment manager by issuing the following command from a command prompt on the MVS console:

```
START dmgr_proc_name,JOBNAME=server_short_name,
 ENV=cell_short_name.node_short_name.server_short_name
```

**Note:** This command must be entered on a single line. It is split here for display purposes.



Then you can launch wsadmin immediately to connect to the deployment manager, or start a node and application server to connect to.

- If an application server is running and you still get this error:
  - If you are running remotely (that is, on a different machine from the one running WebSphere Application Server), you must use the **-host** *hostname* option to the wsadmin command to direct wsadmin to the right physical server.
  - If you are using the **-host** option, try pinging the server machine from the command line from the machine on which you are trying to launch wsadmin to verify there are no issues of connectivity such as firewalls.
  - verify that you are using the right port number to connect to the WebSphere Application Server process:
    - If you are not specifying a port number (using the **-port** option) when you start wsadmin, wsadmin uses the default port specified in *install\_dir/properties/wsadmin.properties*, property `name=com.ibm.ws.scripting.port` (default value =8879).
    - The port that wsadmin should send on depends on the server process wsadmin is trying to connect to.

For a single-server installation, wsadmin attempts to connect to the application server process by default. To verify the port number:

- Look in the file *install\_dir/config/cells/node\_name/nodes/node\_name/serverindex.xml* for a tag containing the property `serverType="APPLICATION_SERVER"`.
- Look for an entry within that tag with the property `endPointName="SOAP_CONNECTOR_ADDRESS"`.
- Look for a **port** property within that tag. This is the port wsadmin should send on.

In a Network Deployment installation, wsadmin launched from the bin directory on the Network Deployment installation attempts to send requests to the deployment manager by default. To verify the port number:

- Get the hostname of the node on which the Deployment Manager is installed.
- Using that hostname, look in *install\_dir/config/cells/node\_nameNetwork/nodes/node\_name Manager/serverindex.xml* for a tag containing the property `serverType="DEPLOYMENT_MANAGER"`.
- Within that tag, look for an entry with a property `endPointName="SOAP_CONNECTOR_ADDRESS"`.
- Within that tag, look for a "port" property. This is the port wsadmin should send on.

**"com.ibm.bsf.BSFException: error while eval'ing Jacl expression: no such method *command name* in class com.ibm.ws.scripting.AdminConfigClient" returned from wsadmin command.**

This error is usually caused by a misspelled command name. Use the **\$AdminConfig help** command to get information about what commands are available. Note that command names are case-sensitive.

**WASX7022E returned from running "wsadmin -c ..." command, indicating invalid command**

If the command following **-c** appears to be valid, the problem may be caused by the shell attempting to do variable substitution. Variable substitution can occur on Unix System Services if wsadmin **-c** invokes a command that is enclosed in double



quotes and includes dollar signs. To confirm that this is the problem, check the command to see if it contains an unescaped dollar sign, for example: `wsadmin -c "$AdminApp install ...."`.

To correct this problem, escape the dollar sign with a backslash. For example: `wsadmin -c "\$AdminApp install ..."`.

**Note:** When the command is enclosed in single quotes, the shell does not attempt to do variable substitution. Therefore, you do not need to escape the dollar sign. Example: `wsadmin.sh -c '$AdminApp install ...'`

**com.ibm.ws.scripting.ScriptingException: WASX7025E: String "" is malformed; cannot create ObjectName**

One possible cause of this error is that an empty string was specified for an object name. This can happen if you use one scripting statement to create an object name and the next statement to use that name, perhaps in an "invoke" or "getAttribute" command, but you don't check to see if the first statement really returned an object name. For example (the following samples use basic Jacl commands in addition to the wsadmin Jacl extensions to make a sample script):

```
#let's misspell "Server"
set serverName [$AdminControl queryNames type=Server,*]
$AdminControl getAttributes $serverName
```

To correct this error, make sure that object name strings have values before using them. For example:

```
set serverName[$AdminControl queryNames type=Server,*]
if {$serverName == ""} {puts "queryNames returned empty - check query argument"}
else {$AdminControl getAttributes $serverName}
```

For details on Jacl syntax beyond wsadmin commands, refer to the Tcl developers' site, <http://www.tcl.tk>.

## Web module or application server dies or hangs

If an application server dies (its process spontaneously closes), or freezes (its Web modules stop responding to new requests):

- Isolate the problem by installing Web modules on different servers, if possible.
- To detect memory leak problems, enable verbose garbage collection on the application server. This feature adds detailed statements to the JVM error log file of the application server about the amount of available and in-use memory. To set up verbose garbage collection:
  1. Select **Servers > Application Servers > server\_name > Process Definition > Java Virtual Machine**, and enable **Verbose Garbage Collection**.
  2. Stop and restart the application server.
  3. Periodically, or after the application server stops, browse the log file for garbage collection statements. Look for statements beginning with "allocation failure". The string indicates that a need for memory allocation has triggered a JVM garbage collection (freeing of unused memory). Allocation failures themselves are normal and not necessarily indicative of a problem. The allocation failure statement is followed by statements showing how many bytes are needed and how many are allocated.

If there is a steady increase in the total amount of free and used memory (the JVM keeps allocating more memory for itself), or if the JVM becomes unable to allocate as much memory as it needs (indicated by the bytes needed statement), there might be a memory leak.

- If the verbose garbage collection output indicates that the application server is running out of memory, one of the following problems might be present:
  - There is a memory leak in application code that you must address. To pinpoint the cause of a memory leak, enable the **RunHProf** function in the Servers > Application Servers > server\_name > Process Definition > Java Virtual Machine pane of the problem application server:
    - In the same JVM pane, set the **HProf Arguments** field to a value similar to `depth=20,file=heapdmp.txt`. This value shows exception stacks to a maximum of 20 levels, and saves the heapdump output to the `install_root/bin/heapdmp.txt` file.
    - Save the settings.
    - Stop and restart the application server.
    - Re-enact the scenario or access the resource that causes the hang or crash, if possible. Stop the application server. If this is not possible, wait until the hang or crash happens again and stop the application server.
    - Examine the file into which the heapdump was saved. For example, examine the `install_root/bin/heapdmp.txt` file:
      - Search for the string, "SITES BEGIN". This finds the location of a list of Java objects in memory, which shows the amount of memory allocated to the objects.
      - The list of Java objects occurs each time there was a memory allocation in the JVM. There is a record of what type of object the memory instantiated and an identifier of a trace stack, listed elsewhere in the dump, that shows the Java method that made the allocation.
      - The list of Java object is in descending order by number of bytes allocated. Depending on the nature of the leak, the problem class should show up near the top of the list, but this is not always the case. Look throughout the list for large amounts of memory or frequent instances of the same class being instantiated. In the latter case, use the ID in the trace stack column to identify allocations occurring repeatedly in the same class and method.
      - Examine the source code indicated in the related trace stacks for the possibility of memory leaks.
  - The default maximum heap size of the application server needs to be increased.
- If an application server spontaneously dies, there will be an SDUMP. See Using the Error Dump and Cleanup interface for instructions on how to analyze the dump.

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

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

## Errors when trying to configure or enable security

What kind of error are you seeing?

- "LTPA password not set validation failed" message displayed as error in the Administrative Console after enabling global security.
- The `setupClient.bat` or `setupClient.sh` file is not working correctly
- "Java HotSpot(TM) Server VM warning: Unexpected Signal 11 occurred under user-defined signal handler 0x7895710a" message occurs in the `native_stdout.log` file when enabling security on the HP-UX11i platform
- If you have successfully configured security (made changes, saved the configuration, and enabled security with no errors), but are now having

problems accessing Web resources or the administrative console, refer to Errors or access problems after enabling security.

### **"LTPA password not set. validation failed" message displayed as error in the Administrative Console after saving global security settings**

This error can be caused if, when configuring WebSphere Application Server security, "LTPA" is selected as the authentication mechanism, and the LTPA password field is not set. To resolve this problem:

- Select Security **Authentication Mechanism** > **LTPA** in the console left-hand navigation pane.
- Complete the password and confirm password fields.
- Click **OK**.
- Try setting Global Security again.

### **The setupClient.bat or setupClient.sh file is not working correctly**

The setupClient.bat file on Windows platforms and the setupClient.sh file on UNIX platforms incorrectly specify the location of the SOAP security properties file.

In the setupClient.bat file, the correct location should be:

```
set CLIENTSOAP=-Dcom.ibm.SOAP.ConfigURL=file:%WAS_HOME%/properties/soap.client.props
```

In the setupClient.sh file, the CLIENTSOAP variable should be:

```
CLIENTSOAP=-Dcom.ibm.SOAP.ConfigURL=file:$WAS_HOME/properties/soap.client.props
```

In the setupClient.bat and setupClient.sh files, complete the following steps:

1. Remove the leading / after file:.
2. Change sas to soap.

### **Java HotSpot(TM) Server VM warning: Unexpected Signal 11 occurred under user-defined signal handler 0x7895710a message occurs in the native\_stdout.log file when enabling security on the HP-UX11i platform**

After you enable security on HP-UX 11i platforms, the following error in the native\_stdout.log file occurs, along with a core dump and WebSphere Application Server does not start:

```
Java HotSpot(TM) Server VM warning:
Unexpected Signal 11 occurred under user-defined signal handler 0x7895710a
```

To work around this error, apply the fixes recommended by HP for Java at the following URL:

<http://www.hp.com/products1/unix/java/infolibrary/patches.html>.

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

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

## **Workload not getting distributed**

What kind of problem are you seeing?

- Web (HTTP) requests are not distributed to all servers.
- Enterprise bean requests are not distributed to all servers.

- A failing server still receives enterprise bean requests (failover fails).

If none of these problem solution descriptions fixes your problem:

1. Browse the JVM logs of the problem deployment manager and application servers:
  - a. Look up any error messages by selecting the **Reference** view of the information center navigation and expanding **Messages** in the navigation tree.
  - b. If Java exceptions appear in the log files, try to determine the actual subcomponent directly involved in the problem by examining the trace stack and looking for a WebSphere Application Server-related class near the top of the stack (names beginning with `com.ibm.websphere` or `com.ibm.ws`) that threw the exception. If appropriate, review the steps for troubleshooting the appropriate subcomponent under the Troubleshooting by component: what is not working? topic.  
For example, if the exception appears to have been thrown by a class in the `com.ibm.websphere.naming` package, review the Naming Services Component troubleshooting tips topic.
2. Ensure that all the machines in your configuration have TCP/IP connectivity to each other by running the **ping** command:
  - a. From each physical server to the Deployment Manager
  - b. From the Deployment Manager to each physical server
3. Although the problem is happening in a clustered environment, the actual cause might be only indirectly related, or unrelated, to clustering. Investigate all relevant possibilities:
  - a. If an enterprise bean on one or more servers is not serving requests, review the Cannot access an enterprise bean from a servlet, JSP, stand-alone program, or other client and Cannot access an object hosted by WebSphere Application Server from a servlet, JSP file, or other client topics.
  - b. If problems seem to appear after enabling security, review the Errors or access problems after enabling security topic.
  - c. If an application server stops responding to requests, or spontaneously dies (its process closes), review the Web module or application server dies or hangs topic.
  - d. If SOAP requests are not being served by some or all servers, review the Errors returned to client trying to send a SOAP request topic.
4. Check to see if the problem is identified and documented by looking at available online support (hints and tips, technotes, and fixes).

### **Web (HTTP) requests are not distributed to all servers**

If HTTP requests are not being distributed to all servers:

- Check your PrimaryServers list. The plug-in load balances across all servers that are defined in the PrimaryServers list, if affinity has not been established. If you do not have a PrimaryServers list defined, the plug-in load balances across all servers defined in the cluster, if affinity has not been established. In the case where affinity has been established, the plug-in should go directly to that server, for all requests within the same HTTP session.
- If some servers are servicing requests and one or more others are not, try accessing a problem server directly to verify that it works, apart from workload management issues. If that does not work:
  - Use the administrative console to ensure that the affected server is running.
  - See the article Web resource (JSP, servlet, html file, image, etc) will not display for more information.
- See the article HTTP plug-in component troubleshooting tips for more information.

## EJB requests are not distributed to all servers

If a client cannot reach a server in a cluster thought to be reachable, a server might be marked unusable, or is down. To verify this:

- Use the administrative console to verify that the server is started. Try starting it, or if started, stop and restart it.
- Browse the administrative console and verify that the node that runs the server having the problem appears. If it does not:
  - Review the steps for adding a node to a cluster.
  - Review the steps in the One or more nodes do not show up in the administrative console topic.
- If possible, try accessing the enterprise bean directly on the problem server to see if there is a problem with TCP/IP connectivity, application server health, or other problem not related to workload management. If this fails, review the topic Cannot access enterprise bean from a servlet, JSP, stand-alone program , or other client.

## A failing server still receives enterprise bean requests (failover fails)

Some possible causes of this problem are:

- The client might have been in a transaction with an enterprise bean on the server that went down. Check the JVM logs of the application server hosting the problem enterprise bean instance. If a request is returned with **CORBA SystemException COMM\_FAILURE org.omg.CORBA.completion\_status.COMPLETED\_MAYBE**, this might be working as designed. The design is to let this particular exception flow back to the client, since the transaction might have completed. Failing over this request to another server could result in this request being serviced twice.
- If the requests sent to the servers come back to the client with any other exceptions consistently, it might be that no servers are available.

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

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

## Errors setting up multiserver environments

What kind of problem are you seeing?

- After creating and starting a cluster, the cluster does not start, and logs show that servers in the cluster are not found.
- One or more nodes do not show up in the administrative console.
- The addNode command fails.
- Application files are not present on all nodes.
- After downloading the Network Deployment plug-in to my system, my server does not start.
- In a clustered environment, a server with debug mode enabled does not start.

If none of these problem solution descriptions fixes your problem:

1. Browse the logs of the problem deployment manager and application servers:
  - a. Look up any error messages by selecting the **Reference** view of the information center navigation and expanding **Messages** in the navigation tree.
  - b. If Java exceptions appear in the log files, try to determine the actual subcomponent directly involved in the problem by examining the trace

stack and looking for a WebSphere Application Server-related class near the top of the stack (names beginning with `com.ibm.websphere` or `com.ibm.ws`) that threw the exception. If appropriate, review the steps for troubleshooting the appropriate subcomponent in the Troubleshooting by component: what is not working? topic.

For example, if the exception appears to have been thrown by a class in the `com.ibm.websphere.naming` package, review the Naming Services Component troubleshooting tips topic.

2. Ensure that all the machines in your configuration have TCP/IP connectivity to each other by running the **ping** command:
  - a. From each physical server to the Deployment Manager
  - b. From the Deployment Manager to each physical server
3. Although the problem is happening in a clustered environment, the actual cause might be only indirectly related, or unrelated, to clustering. Investigate all relevant possibilities:
  - a. If an enterprise bean on one or more servers is not serving requests, review the Cannot access an enterprise bean from a servlet, JSP, stand-alone program, or other client and Cannot access an object hosted by WebSphere Application Server from a servlet, JSP file, or other client topics.
  - b. If problems seem to appear after enabling security, review the Errors or access problems after enabling security topic.
  - c. If an application server stops responding to requests, or spontaneously dies (its process closes), review the Web module or application server dies or hangs topic.
  - d. If SOAP requests are not being served by some or all servers, review the Errors returned to client trying to send a SOAP request topic.
4. Check to see if the problem is identified and documented by looking at available online support (hints and tips, technotes, and fixes).

#### **After creating and starting a cluster, the cluster does not start, and logs show that servers in the cluster are not found**

This error can occur when the configuration is not synchronized from the deployment manager to a node. If *auto synchronization* is enabled, wait until the synchronization has had a chance to run. If you are using manual synchronization, explicitly request a sync to each node on the cluster.

To determine whether synchronization has taken place, look at the configuration on the node machines using the administrative console and verify that the new cluster members are defined on each node.

#### **One or more nodes do not show up in the administrative console**

This can occur when there is a basic connectivity problem between the deployment manager server and other servers in the topology. To determine whether this is the problem, look for the `fileserversindex.xml` in the deployment manager directory structure.

- If the problem node does not appear in the list, review the steps for adding a node to the cluster.
- If the problem node does appear in the list:
  - From the deployment manager server, ping the server name as it appears in the list. If the ping command shows no communication, verify that the hostname is correct in the list, and correct it if necessary, then restart the deployment manager.



- If the name that appears in the list is the short name, ping the fully qualified network name. If the corrected name works, update the list and restart the deployment manager.
- If the problem server uses Dynamic Host Configuration Protocol (DHCP), try replacing the logical hostname with the IP address and restart the deployment manager. If this resolves the problem, be aware that you must change `serverindex.xml` each time the problem server address changes, potentially each time the problem machine is rebooted. To avoid this problem, consider assigning a static IP address to the server.
- If you still cannot establish communication between the servers, contact your network administrator to resolve the problem, and restart the deployment manager after the problem is corrected.

### The `addNode` command fails

This error can occur when the deployment manager Domain Name Server (DNS) configuration is set up improperly. The default installation on Linux uses the loopback address (127.0.0.1) as the default host address. To verify that this is the problem, query the hostname of the suspect machine. If it returns localhost 127.0.0.1, or if file transfer traces at the node show the node trying to upload files to a URL that includes 127.0.0.1, the node has an incorrect DNS configuration.

To correct this problem, update the `/etc/hosts` file or the name service configuration file, `/etc/nsswitch.conf`, to query the Domain Name Server or Network Information Server (NIS) before searching hosts.

### Application files are not present on all nodes

In the WebSphere Application Server Network Deployment environment, application binary files are transferred to the individual nodes where applications are supported as part of the **node sync** operation. During node sync, application files are only propagated if their deployment descriptors specify **enableDistribution=true**. This flag is specified as part of the application installation procedure in the administrative console, and is stored as a property in the `install_root/config/cells/cell_name/applications/application_name/deployment.xml` file.

To confirm that this is the cause, check to see whether the `enableDistribution` flag is set. If it is already set to true, ensure that the target node is configured to run auto file synchronization.

If both of these settings are correct and the problem persists, manually perform an explicit synchronization. If the application files still do not appear in the installation directory, use the `EARExpander` tool (located in `install_root/bin`) to expand the EAR file from the repository to the installation destination. On remote nodes, the repository should appear in `theconfig/cells/cell_name/applications/application_name.ear/` directory.

### After downloading the Network Deployment plug-in to my system, my server does not start

If you experience this situation, the most likely cause is that the transport paths in the plug-in must be modified to work in your environment. See the Example: Manually editing transport settings in the `server.xml` file topic for information on how to modify these settings.



## In a clustered environment, a server with debug mode enabled does not start

This problem occurs when the following three conditions exist:

- Multiple server processes are configured to run on the same node
- More than one server has Debug Mode enabled
- The debug arguments for more than one of the servers have been left at the default values, so that more than one server in the node is trying to use the same debug port (port number 7777).

The server will not start because multiple servers processes running on the same physical host machine with debug enabled cannot use the same debug port.

To correct this problem, for each server:

1. On the Administrative Console select **Server > Application Servers > *server name* > Process Definition > *process name* > Java Virtual Machine**
2. Update the Debug argument so that the address of the debug port (address=*port number*) is unique for each server process.

## Cannot uninstall an application or remove a node or application server

What kind of problem are you having?

- After uninstalling an application through wsadmin tool, the application continues to run and throws "DocumentIOException"
- The removeNode command does not remove the installed application from the deployment manager
- I cannot display the syntax for the removeNode command.

If none of these steps fixes your problem:

- Make sure that the application and its Web and EJB modules, are in a stopped state before uninstalling.
- If you are uninstalling or installing an application using **wsadmin**, make sure that you are using the **-conntype NONE** option to invoke **wsadmin** and enable local mode. To use the **-conntype NONE** option, stop the hosting application server before uninstalling the application.
- Check to see if the problem has been identified and documented by looking at the available online support (hints and tips, technotes, and fixes).

### After uninstalling application through the wsadmin tool, the application throws "DocumentIOException"

If this exception occurs after the application was uninstalled using wsadmin with the **-conntype NONE** option:

- Restart the server or,
- Rerun the uninstall command without the **-conntype NONE** option.

### The removeNode command does not remove the installed application from the deployment manager

If the applications were installed indirectly using the **addNode** program with the **-includeapps** option, then **removeNode** will not uninstall them, since they may be in use by other nodes. These applications must be explicitly uninstalled, for example through the administrative console.

## I cannot display the syntax for the removeNode command

Unlike the addNode command, the removeNode command is valid with no parameters, so executing it will execute the operation, that is, remove the node, without displaying the command syntax.

To see the valid options for removeNode, execute removeNode -? or removeNode -help.

## Problems creating or using HTTP sessions

**Note:** To view and update the Session Manager settings discussed here, use the administrative console. Select the application server that hosts the problem application, then under **Additional properties**, select **Web Container**, then **Session manager**.

What kind of problem are you having?

- HTTP Sessions are not getting created, or are lost between requests.
- HTTP Sessions are not persistent (session data lost when application server restarts, or not shared across cluster).
- Session is shared across multiple browsers on same client machine.
- Session is not getting invalidated immediately after specified Session timeout interval.
- Unwanted sessions are being created by jsps.

If your problem is not described here, or none of these steps fixes the problem:

•

### HTTP Sessions are not getting created, or are lost between requests

By default, the Session Manager uses cookies to store the session ID on the client between requests. Unless you intend to avoid cookie-based session tracking, ensure that cookies are flowing between WebSphere Application Server and the browser:

- Make sure the **Enable cookies** checkbox is checked under the **Session tracking Mechanism** property.
- Make sure cookies are enabled on the browser you are testing from or from which your users are accessing the application.
- Check the Cookie domain specified on the SessionManager (to view the or update the cookie settings, in the **Session tracking mechanism->enable cookies** property, click **Modify**).
  - For example, if the cookie domain is set as ".myCom.com", resources should be accessed using that domain name, e.g. `http://www.myCom.com/myapp/servlet/sessionservlet`.
  - If the domain property is set, make sure it begins with a dot (.). Certain versions of Netscape do not accept cookies if domain name doesn't start with a dot. Internet Explorer honors the domain with or without a dot. For example, if the domain name is set to `mycom.com`, change it to `.mycom.com` so that both Netscape and Internet Explorer honor the cookie.
- Check the **Cookie path** specified on the SessionManager. Check whether the problem url is hierarchially below the Cookie path specified. If not correct the Cookie path.
- If the Cookie maximum age property is set, ensure that the client (browser) machine's date and time is the same as the server's, including the time zone. If the client and the server time difference is over the "Cookie maximum age" then every access would be a new session, since the cookie will "expire" after the access.

- If you have multiple web modules within an enterprise application that track sessions:
  - If you want to have different session settings among web modules in an enterprise application, ensure that each web module specifies a different cookie name or path, or
  - If Web modules within an enterprise application use a common cookie name and path, ensure that the HTTP session settings, such as Cookie maximum age, are the same for all Web modules. Otherwise cookie behavior will be unpredictable, and will depend upon which application creates the session. Note that this does not affect session data, which is maintained separately by Web module.
- Check the cookie flow between browser and server:
  1. On the browser, enable "cookie prompt". Hit the servlet and make sure cookie is being prompted.
  2. Access the session servlet from the browser.
  3. The browser will prompt for the cookie; note the jsessionid.
  4. Reload the servlet, note down the cookie if a new cookie is sent.
  5. Check the session trace and look for the session id and trace the request by the thread. Verify that the session is stable across web requests:
    - Look for **getHttpSession(...)** which is start of session request.
    - Look for **releaseSession(..)** which is end of servlet request.
- If you are using URL rewriting instead of cookies:
  - Ensure there are no static HTML pages on your application's navigation path.
- If you are using SSL as your session tracking mechanism:
  - Ensure that you have SSL enabled on your IBM HTTP Server or iPlanet HTTP server.
- If you are in a clustered (multiple node) environment, ensure that you have session persistence enabled.

### HTTP Sessions are not persistent

If your HTTP sessions are not persistent, that is session data is lost when the application server restarts or is not shared across the cluster:

- Check the Datasource.
- Check the SessionManager's Persistence Settings properties:
  - If you intend to take advantage of Session Persistence, verify that Persistence is set to **Database**.
  - If you are using **Database-based persistence**:
    - Check the jndi name of the datasource specified correctly on SessionManager.
    - Specify correct userid and password for accessing the database.

Note that these settings have to be checked against the properties of an existing Data Source in the admin console. The Session Manager does not automatically create a session database for you.

    - The Datasource should be non-JTA, i.e. non XA enabled.
    - Check the logs for appropriate database error messages.
    - With DB2, for row sizes other than 4k make sure specified row size matches the DB2 page size. Make sure tablespace name is specified correctly.

### Session is shared across multiple browsers on same client machine

This behavior is browser-dependent. It varies between browser vendors, and also may change according to whether a browser is launched as a new process or as a subprocess of an existing browser session (for example by hitting Ctl-N on Windows).

The Cookie maximum age property of the Session Manager also affects this behavior, if cookies are used as the session-tracking mechanism. If the maximum age is set to some positive value, all browser instances share the cookies, which are persisted to file on the client for the specified maximum age time.

### **Session is not getting invalidated immediately after specified Session timeout interval**

The SessionManager invalidation process thread runs every  $x$  seconds to invalidate any invalid sessions, where  $x$  is determined based on the Session timeout interval specified in the Session manager properties. For the default value of 30 minutes,  $x$  is around 300 seconds. In this case, it could take up to 5 minutes (300 seconds) beyond the timeout threshold of 30 minutes for a particular session to become invalidated.

### **Unwanted sessions are being created by jsps**

As required by the JavaServer Page specification, jsps by default perform a `request.getSession(true)`, so that a session is created if none exists for the client. To prevent jsps from creating a new session, set the session scope to **false** in the jsp file using the page directive as follows:

```
<% @page session="false" %>
```

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

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

## **JSP source code shown by the Web server**

### **Problem**

If you share the document root of the WebSphere Application Server within the Web server document root, a security exposure can result as the Web server might display the JSP source file as plain text.

You can use the WebSphere Web server plug-in set of rules to determine whether a given request will be handled by the WebSphere Application Server. When an incoming request fails to match those rules, the Web server plug-in returns control to the Web server so that the Web server can fulfill the request. In this case, the unknown host header causes the Web server plug-in to return control to the Web server because the rules do not indicate that the WebSphere Application Server should handle it. Therefore, the Web server looks for the request in the Web server document root. Since the JSP source file is stored in the document root of the Web server, the Web server finds the file and displays it as plain text.

### **Suggested solution**

Move the WebSphere Application Server JSP source file outside of the Web server document root. Then, when this request comes in with the unknown host header, the plug-in returns control to the Web server and the JSP source file is not found in the document root. Therefore, the Web server returns a 404 File Not Found error rather than the JSP source file.

## Problems using tracing, logging or other troubleshooting features

What kind of problem are you having?

- Netscape browser fails when trying to enable a component trace

### Netscape browser fails when trying to enable a component trace

On systems using AIX, the Netscape browser fails when you try to enable trace on a component.

To work around this problem, do one of the following:

- Disable JavaScript on the browser and continue setting trace.
- Administer the AIX server from a remote machine running another browser and operating system.
- Change the trace manually in the *server.xml* file.

## Errors connecting to the administrative console from a Netscape browser

What kind of problem are you having?

- Resizing the Netscape browser results in an error (nresize)
- Resizing the Netscape browser causes an error 404 message (404)
- Netscape screen blanks out while using the administrative console (152339)
- Resizing Netscape Version 4.7 causes errors (resizing47)
- Enabling Netscape Version 4.7 to display double-byte character set correctly (netscape479)
- Using Netscape Version 4.79 on a Solaris Operating Environment causes problems (479)
- Limitations occur when using Netscape with Solaris Operating Environment (nhas)
- Netscape browser must be capable of launching from a terminal window on all UNIX platforms (151829)

If you are able to bring up the browser page, but the console's behavior is inconsistent, error-prone, or unresponsive, try upgrading the browser you are using. Older browsers may not support the administrative console's features. For a listing of supported Web browsers, see <http://www.ibm.com/software/webservers/appserv/doc/latest/prereq.html>.

### Resizing the Netscape browser results in an error (nresize)

If you resize your Netscape browser, you could get a "Data Missing" error. The error message disappears in 60 seconds.

### Resizing the Netscape browser causes an error 404 message (404)

When connecting to the IBM WebSphere Application Server administrative console from a Netscape browser, resizing the browser can cause an error 404 message to occur. This situation occurs because the browser reloads the frame when resizing the window.

To avoid getting the error message, you can refrain from resizing the Netscape browser window, or you can connect to the IBM WebSphere Application Server administrative console using an Internet Explorer browser.

## **Netscape screen blanks out while using the administrative console (152339)**

While working with the right-hand panel of the administrative console to do regular administrative tasks, the browser screen blanks out intermittently.

To work around this problem, do one of the following:

- After the problem occurs, close the Netscape browser, log in again, and continue working.
- Use the Internet Explorer browser from a Windows machine.
- Use Netscape 7.x, Mozilla 1.x, Opera 5, or Konquerer browsers on the platform, depending on which is available. Although there is not formal support for these browsers, they have all been used successfully with the product and in many cases work better than the previous 4.7.x series of Netscape browsers

## **Resizing Netscape Version 4.7 causes errors (resizing47)**

You receive the following error messages when resizing Netscape Version 4.7:

```
Error 0
 An error occurred while processing request: http://localhost:9090/admin/upload.do
 message:
 Details
 com.ibm.websphere.servlet.error.ServletErrorReport:
 at java.lang.Class.newInstance0(Native Method)
 ...
```

After resizing Netscape 4.7, Netscape has to reload the page just as it initially loads the page on the first request. For pages that do not expect POST data, it is not a problem. But for pages that do, Netscape 4.7 cannot retain the data.

## **Enabling Netscape Version 4.7 to display double-byte character set correctly (netscape479)**

When using Netscape on AIX platforms without the translated package, the English version of Netscape is available for all locale environments as the default package. However, the English version of Netscape does not display double-byte character set (DBCS) characters on the browser radio buttons and title bars because the fonts are mismatched.

To work around this problem, you can install the message resource to make the translated version of Netscape available on a DBCS environment. Use the translated version of Netscape to display the corrupted DBCS correctly. Change the locale from English to the expected DBCS locale before starting Netscape. For example, issue the following commands to display the Japanese contents on Ja\_JP (AIX Shift JIS locale):

```
$ export LANG=Ja_JP
$ netscape&
```

## **How to check the installed package**

For the AIX platform:

1. Run the `lspp -aL | grep Netscape.msg` command.
2. Check the language and locale information using Netscape.msg package. For example, the result of the command run in step 1 is as follows for the system having Japanese-translated Netscape:

```
ls1pp -aL | grep Netscape.msg
Netscape.msg.Ja_JP.communicator.rte
Netscape.msg.ja_JP.communicator.rte
```

For the Red Hat Linux Version 7.2 platform:

Run the **rpm-ql netscape-common** command to get the list of installed files.

In case of Netscape Version 4.78 on Red Hat Linux Version 7.2, the translated packages for seven languages including four DBCS languages (Japanese, Korean, Simplified Chinese, and Traditional Chinese) are included in a common package.

### How to install the package

For the AIX platform:

Install the Netscape.msg.%L.communicator.rate package where %L is the locale information for the System Management Interface Tool (SMIT) or SMITTY to use. This message resource is provided by the Bonus Pack CD or the Expansion Pack CD for AIX platform, which provides the Netscape common package.

Perform the following steps to install the package:

1. Set the media which has the Netscape installation image. For AIX 5L platform, set the Expansion Pack media.
2. Go to the directory which has the Netscape installation package. For AIX 5L platform Expansion Pack, go to the installp/ppc directory.
3. Start SMIT or SMITTY.
4. Click System Environments > Manage Language Environment > Change/Show Applications for a Language > Add Language for Application Already Installed.
5. In the Add Language for Application Already Installed field select a language (locale) from the drop-down menu list for the Language translation to install field.
6. Click OK.
7. In the Add Language for Application Already Installed field enter the directory location. Enter "." for the current directory or /cdrom/installp/ppc for the full path as input directory for the software. Do not select cdrom as the input device because the installation package does not exist on the root directory of the CD-ROM.
8. Click OK.
9. Click List to check if the configuration is correct. The configuration is correct if Netscape is on the list of installed application. Start the installation now.
10. Click OK.
11. Exit SMIT or SMITTY.

For Red Hat Linux Version 7.2 platform:

The default Netscape installation can install each translated resource as part of the common package. There is no way to install the translated package individually. However, TrueType fonts (TTF packages) require installation for each language to correctly display the translated resources on Netscape. Use the following to check if the required font files are correctly installed:



```
$ rpm -qa | grep ttf
resources.0-4
ttfm-0.9.1-8
ttfonts-ja-1.0-6
ttfonts-ko-1.0-10
ttfonts-zh_CN-2.11-5
ttfonts-zh_TW-2.11-5
```

### **Using Netscape Version 4.79 on a Solaris Operating Environment causes problems (479)**

Using Netscape Version 4.79 on a Solaris Operating Environment to access the administrative console causes problems with some key text translations with the zh\_TW.EUC locale. This situation is not a problem when you use Netscape Version 4.7. The officially supported version of Netscape on a Solaris Operating Environment is Version 4.79, but in this case the workaround is to use Netscape Version 4.7.

### **Limitations occur when using Netscape with Solaris Operating Environment (nhas)**

If you click Troubleshooting > Logs and Trace > *server* > Diagnostic Trace > Modify, the window that pops up allowing you to select the Components and Groups to trace might not display a scroll bar, preventing you from viewing all the components and groups.

The text area displaying the selected components, groups, and trace levels does not have a vertical scroll bar. This omission is a limitation of Netscape on a Solaris Operating Environment.

To work around this problem, refresh the window to show the scroll bar.

### **Netscape browser must be capable of launching from a terminal window on all UNIX platforms (151829)**

To make sure the Netscape browser can launch from a terminal window, edit or create a file called `profile` in the `/etc` directory and add the Netscape directory to the system path. For example: `PATH=$PATH:/opt/Netscape export PATH`

---

## **Resolving timeout conditions**

This file gives an overview of how to resolve timeout conditions

In such a complex environment as WebSphere Application Server for z/OS, timeouts might occur for many different reasons. Although you can alter timeout values, you should not do so until you understand why the timeout occurs. Depending on the timeout condition, you might be able to permanently fix the timeout condition by doing some system or application tuning. For example, if the diagnostic data indicates throughput problems, you can alter the number of server regions, the number of threads within each server region, or the use of replicated servers.

Generally speaking, increasing the timeout values should be your last resort, or only a temporary action taken to prevent multiple timeout-abend dumps from causing system performance problems. If you increase timeout values without

properly diagnosing the timeout condition, the only results you might see are less frequent abends and dumps for the same timeout condition, or slower system or application performance.

## Understanding how timers work

This file gives an overview of understanding timers

Timers define a limit to the amount of time required for a specific operation to complete. When the timer begins its countdown depends on type of operation it controls. The timers that WebSphere Application Server for z/OS uses can be classified into the general types described later in this article; the specific timers themselves are described in Controlling behavior through timeout values. Most of the timers have a default value that defines a reasonable limit for the particular operation to complete. When the timer pops (that is, reaches the time limit), WebSphere Application Server for z/OS takes one of the following actions:

- Sends a minor code to the client for timers that pop before the client request is dispatched to a servant region.
- Abnormally ends the servant region with an EC3 ABEND for timers that pop while the client request is being processed by an application component running in the servant region. All threads in the abending servant region will be terminated.

WebSphere terminates the servant region to prevent the application from tying up resources, thus causing other requests to start backing up. Once the servant is terminated, WLM starts a new servant to take its place and continue processing requests from the controller.

Different types of timers might be counting down simultaneously, because the operations they control might overlap to a certain degree. For example, suppose the application server receives an IIOP client request that will be processed by an application component that uses transaction support. In this case, both of the following WebSphere timers can be counting down simultaneously:

- `control_region_wlm_dispatch_timeout`, which limits both the amount of time a client request waits on the WLM queue, as well as the time required for the application component to process the request; and
- `transaction_defaultTimeout`, which limits the amount of time the controller will wait for a transaction to be either committed or rolled back.

These timers overlap only for the time during which the application's transaction is being processed. To determine which timer cause the error, you can use the symptoms- specific minor codes or EC3 ABEND reason codes.

### General types of timers and the operations they control

| General type | Timer processing                                                                                                                                                                                                           | Timeout symptoms       |
|--------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------|
| <b>Input</b> | Input timers define limits for receiving a complete request; the countdown starts when a connection to the J2EE server is established. The communication protocol (HTTP, HTTPS) determines the timer used for the request. | The session is closed. |

| General type        | Timer processing                                                                                                                                                                                                                                                                                                                                                                 | Timeout symptoms                                                                                                                                                             |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Session</b>      | Session timers define limits for the use of session connections. These timers start the countdown as soon as a session becomes idle.                                                                                                                                                                                                                                             | The session is closed.                                                                                                                                                       |
| <b>WLM dispatch</b> | Dispatch timers control how long a complete client request waits to be dispatched in a servant region for processing. The countdown starts when the controller places the request on the WLM queue. Depending on the specific timer, the time limit can include not only wait time on the WLM queue, but also the time required for processing a response to the client request. | Message BBOO0232W and an EC3 ABEND in the servant (region), with one of these accompanying reason codes:                                                                     |
| <b>Transaction</b>  | Transaction timers define how long: <ul style="list-style-type: none"> <li>An application or controller will wait for one transaction to complete. The countdown starts when the container starts a transaction on behalf of the application component.</li> <li>A controller will attempt to recover in-doubt transactions during peer restart and recovery mode.</li> </ul>    | <ul style="list-style-type: none"> <li>Message BBOT0003W or BBOO0232W</li> <li>An EC3 ABEND in the servant (region), with one of these accompanying reason codes:</li> </ul> |
| <b>Output</b>       | Output timers define how long a controller will wait to receive output for a client request. The countdown starts when the client request is dispatched to the servant region for processing. The communication protocol (HTTP, HTTPS) determines the timer used for the request.                                                                                                | Message BBOO0232W and an EC3 ABEND in the servant (region), with reason code 04130007                                                                                        |

## Guidelines for analyzing diagnostic data for timeout conditions

This file gives an overview of how to enable and use the System Management Facilities (SMF) to collect and record system and job-related information.

The following guidelines provide instructions for finding diagnostic data in an SVC dump that can help you determine what timeout condition occurred:

- Find the task with the EC3 abend:

1. Format the TCB summary for the servant that was timed out by entering the following command:

```
ip summ format asid(x' address ')
```

where *address* is the address space ID of the servant.

Find the TCB that had the EC3 completion code. Ignore any EC3 completion code on the "main" thread which is the 4th TCB listed in the summary format (the 1st one after the 3 MVS TCBs). The WebSphere main thread is the one that is waiting in BBO\_BOA::impl\_is\_ready. No application requests are ever dispatched on this thread, therefore there is nothing to timeout. During timeout processing the main thread for the server region is also abended with EC3 as a mechanism of bringing the address space down. Thus the reason why the EC3 completion code may appear on the main thread. This is never the cause of a timeout though, only a result of timeout processing.

2. If there is no EC3 completion code in the TCB summary, look in systrace. Format the systrace in gmt time since the other timestamps you'll be comparing it to are in gmt time. To format in gmt time, enter the following command::

```
ip systrace all time(gmt)
```

You may not see the EC3 abend in systrace either as systrace can cover a small amount of time.

3. You can also try looking in ip verbx mtrace or in syslog to see when the EC3 abend occurred. You'll need this time to determine the 'end' time of the request which is the gmt time the timeout value was reached.
- Determine what timeout values are in effect by checking the reason code associated with the EC3 abend.

| Reason code | Explanation                                                                                                                                                                                                                                                                                                                                    |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 04130002    | The controller issued an ABTERM for this servant region because a transaction timeout occurred. Code under dispatch could have been in a tight loop.                                                                                                                                                                                           |
| 04130003    | The controller issued an ABTERM for this servant region because it was hung trying to move a controller request into the servant region. The target request was timed out, but the servant was currently copying the request. The controller checked the servant for progress at regular intervals, before taking action by issuing an ABTERM. |
| 04130004    | The controller issued a ABTERM for this servant region because the WLM queue timeout occurred. Code under dispatch could have been in a tight loop.                                                                                                                                                                                            |
| 04130005    | The controller issued an ABTERM for this servant region because a transaction timeout occurred. The transaction has timed out, but no current request associated with the transaction was found. The servant associated with the transaction will be terminated.                                                                               |

| Reason code | Explanation                                                                                                                                                                                                                                    |
|-------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| 04130006    | A controller thread encountered a problem while processing a request. The request has been queued to WLM and associated with a servant region. The termination of the associated servant region is needed to complete cleanup for the request. |
| 04130007    | The controller issued a ABTERM for this servant region because the HTTP OUTPUT timeout occurred. Code under dispatch could have been in a tight loop.                                                                                          |

- Find the method name to determine if it was

`httpRequest`

,

`httpsRequest`

or

`DispatchbyURI`

or some other method.

If the request is not specifically a request that came through the HTTP or HTTPS transport handlers, the

`protocol_http_output_timeout`

(HTTP) and

`protocol_https_timeout_output`

(HTTPS) timeout values will not be a factor. In other words, when the request is a

`DispatchbyURI`

method, the request is received through the RMI/IIOP protocol, so the

`protocol_http`

\* variables have no affect.

- Obtain the callback stack for the request, using the IPCS verbexit LEDATA, with the CEEDUMP or NTHREADS option.

## Identifying possible causes of and fixes for timeout conditions

This file lists common timer variables and tools for monitoring these timeout conditions

The timer that expires first might not indicate the actual problem that needs to be fixed. To properly diagnose timeout conditions, you should know all of the timer values that might be in effect for a particular servant region.

| General type of timer | Possible causes                                                            | Possible solutions                                                                                                                     |
|-----------------------|----------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------|
| Input                 | The client sent only part of the data and was delayed in sending the rest. | The application on the client side may want to consider having retry logic in place if it does receive a timeout minor code in return. |

| General type of timer | Possible causes                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           | Possible solutions                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Session               | The session is idle through lack of use.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  | If you consider losing idle sessions to be a problem, increase the values of the persistent-session timeouts, or use the session more frequently.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| WLM dispatch          | <p>No threads are free to pick up the request because of one of the following conditions:</p> <ul style="list-style-type: none"> <li>• The threads are all busy processing requests.</li> <li>• The currently executing threads are waiting for a response from DB2, WebSphere MQ, another server, and so on. In this case, look for messages indicating contention for resources; for example, on the z/OS console, you might see messages about DB2 deadlocks.</li> </ul> <p>In either case, the request times out waiting in the WLM queue to be dispatched in a servant (region).</p> | <p>The case where the threads are all busy processing requests could indicate one of the following conditions:</p> <ul style="list-style-type: none"> <li>• The number of servant regions that WLM may start is set too low. To set this value, on the administrative console select <b>Servers &gt;Application Servers &gt;server_name &gt;Server Instance</b>. Click on <b>Multiple Instances Enabled</b> and specify a value for <b>Maximum Number of Instances</b>.</li> <li>• The number of threads allowed within a servant region is set too low. The number is controlled by the Isolation Policy setting in Administrative console or WebSphere variable: <code>server_region_workload_profile</code></li> <li>• You need to replicate servers to handle the amount of incoming work.</li> </ul> <p>All of these conditions indicate that performance tuning might be necessary.</p> |
| Transaction           | <p>Possible causes of transaction timeouts include:</p> <ul style="list-style-type: none"> <li>• The same as those for WLM dispatch timeouts, or</li> <li>• Delays that prevent the transaction coordinator from committing or rolling back a transaction within the allotted time.</li> </ul>                                                                                                                                                                                                                                                                                            | See the possible solutions for WLM dispatch timeouts. In addition, you can look for messages indicating contention for resources that are involved in the transaction that timed out.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
| Output                | Possible causes of output timeouts are the same as those for WLM dispatch (dispatch is for IIOP, output is for HTTP).                                                                                                                                                                                                                                                                                                                                                                                                                                                                     | See the possible solutions for WLM dispatch timeouts. In addition, you can use the WebSphere variable <code>protocol_accept_http_work_after_min_srs=1</code> to prevent the HTTP transport handler from dispatching requests until WLM starts a minimum number of servant regions.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |

## Common timeout settings

Use this page to configure timeout settings.

This file lists common timer variables. Generally speaking, increasing the timeout values should be your last resort, or only a temporary action taken to prevent multiple timeout-abend dumps from causing system performance problems.

### **control\_region\_jms\_request\_timeout**

Specifies the time, in seconds, that the Controller Region waits for a JMS request to complete. Requests include starting or stopping the Queue Manager, Channel Initiator and Listener.

Specifies the time, in seconds, that the Controller Region waits for a JMS request to complete. Requests include starting or stopping the Queue Manager, Channel Initiator and Listener.

|                  |              |
|------------------|--------------|
| <b>Data Type</b> | Integer      |
| <b>Units</b>     | Seconds      |
| <b>Default</b>   | 60           |
| <b>Range</b>     | 0 to 2147040 |

## Guidelines for altering timeout values

This file lists common timer variables and tools for monitoring these timeout conditions

Generally speaking, increasing the timeout values should be your last resort, or only a temporary action taken to prevent multiple timeout-abend dumps from causing system performance problems. If you increase timeout values without properly diagnosing the timeout condition, the only results you might see are less frequent abends and dumps for the same timeout condition, or slower system or application performance.

For information on how to set values for these timer variables, and how these variables map to internal variables, see Controlling behavior through timeout values

### Common timer variables and tools for monitoring these timeout conditions

| WebSphere variable and its relationship, if any, to other timers                                                                                            | How to monitor processing for this type of timeout condition:                                                                                                                           | If you want to adjust the value, consider the following:                                                                                                                                                                                                                                                         |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>WLM timeout</b><br><br>For HTTP work, the WLM timer is not set and only the ConnectionResponseTimeout is in effect (covering the entire dispatch window) | SMF provides data on WLM queue time                                                                                                                                                     | How long work takes to get to a servant depends on the number of servants that WLM starts, how many you let it start, how many service classes the work is spread across, how much work you're getting, and so on.                                                                                               |
| <b>ConnectionIOTimeout</b><br><br>None.                                                                                                                     | This behavior is not easily monitored. Turning on a trace point would indicate whether a client failed because of this input timeout setting, but tracing has performance consequences. | <ul style="list-style-type: none"> <li>• How long are you willing to allow a control region worker thread to be blocked while it is waiting for a message?</li> <li>• How big are incoming HTTP requests? The larger they are, the longer it might take to get the whole request through the network.</li> </ul> |



| WebSphere variable and its relationship, if any, to other timers                                                                                                                                                                                                                                                                                                                                     | How to monitor processing for this type of timeout condition:                                                                                                         | If you want to adjust the value, consider the following:                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <p><b>ConnectionResponseTimeout</b></p> <p>If the application component starts transactions, then the transaction timers also might be involved.</p>                                                                                                                                                                                                                                                 | <p>This behavior is not easily monitored, but the controller will terminate the servant (region) with abend EC3 for this timeout condition.</p>                       | <ul style="list-style-type: none"> <li>• How long are you willing to let a client hang waiting for a response?</li> <li>• How long are you willing to let a thread in a servant (region) be tied up working on a response before concluding that the request has taken too long?</li> <li>• If you have multiple application threads in the servant (region), all of them will be terminated when only one of them times out. This loss of work might make you want to allow these time outs to occur less frequently.</li> </ul> |
| <p><b>ConnectionKeepAliveTimeout</b></p> <p>None. All the other timers relate to work processing, whereas this one relates to what happens when there is no work.</p>                                                                                                                                                                                                                                | <p>None.</p>                                                                                                                                                          | <p>How much time passes between requests vs. how much does it cost to establish a new session. You would want to keep idle sessions around for a while to avoid the startup cost of a new session, but don't want to keep them forever as resource usage accumulation will eventually be a problem.</p>                                                                                                                                                                                                                           |
| <p><b>Request Timeout (ORB Service)</b></p> <p>None. This variable is a client-side timeout, and IIOF only.</p>                                                                                                                                                                                                                                                                                      | <p>None, other than to observe the timeouts occurring on the client side.</p>                                                                                         | <p>How long are you willing to let the client wait?</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <p><b>ORB listener keep alive ORB SSL listener keep alive</b></p> <p>None. These variables relate to session activity during idle periods and only for IIOF, so these timers do not interact with the ConnectionKeepAliveTimeout timer.</p>                                                                                                                                                          | <p>You should read TCP/IP APAR PQ18618 for information about the SOCK_TCP_KEEPALIVE values and their consequences.</p>                                                | <p>Is it useful to have idle sessions timeout? They normally don't which can consume resources. However, detecting a timeout requires network traffic between TCP/IP stacks. Creating traffic on otherwise idle sessions may have network consequences you don't want.</p>                                                                                                                                                                                                                                                        |
| <p><b>Total Transaction Lifetime Timeout</b></p> <p>This variable can be overridden by applications up to the maximum indicated by the Maximum Transaction Timeout variable, which limits the amount of time an application can set for its transactions to complete. Output timers also might cause work to time out, but the transaction timers and output timers are not aware of each other.</p> | <p>The controller issues message BBOT0003W to indicate a timeout condition, and terminates the servant (region) with abend EC3 reason codes 04130002 or 04130005.</p> | <ul style="list-style-type: none"> <li>• How long are you willing to let a client hang waiting for a response?</li> <li>• How long are you willing to let a thread in a servant (region) be tied up working on a response before concluding that the request has taken too long?</li> <li>• If you have multiple application threads in the servant (region), all of them will be terminated when only one of them times out. This loss of work might make you want to allow these timeouts to occur less frequently.</li> </ul>  |
| <p><b>Maximum Transaction Timeout</b></p> <p>If set, this variable limits the amount of time an application can set for its transactions to complete. If the Maximum Transaction Timeout variable is not set, application transactions are controlled by the time limit set on the Total Transaction Lifetime Timeout variable.</p>                                                                  | <p>None.</p>                                                                                                                                                          | <p>Same considerations as for transaction_defaultTimeout</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <p><b>transaction_recoveryTimeout</b></p> <p>None</p>                                                                                                                                                                                                                                                                                                                                                | <p>None.</p>                                                                                                                                                          | <p>Locks are held while one controller (region) waits for other controllers that are required to resolve in-doubt transactions. How long can you afford to have these resources held?</p>                                                                                                                                                                                                                                                                                                                                         |

## Debugging client exceptions

Start with the client and work your way backward to find the problem. When tracing exceptions back to the original problem, be aware that the RMI/IIOF protocol requires that some exceptions undergo conversion from one type to another as the exception passes through the runtime. Usually this transformation is between CORBA::SystemExceptions and RMI RemoteExceptions. Pay special attention to the CORBA::SystemException minor codes which indicate that a type transformation has occurred.

|                   |                                                                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                             |
|-------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Caused by:</b> | System exception (thrown by runtime)                                                                                                                                                                                                                                                                                                                                                | User exception (thrown by application code)                                                                                                                                                                                                                                                 |
| <b>Look for:</b>  | <ul style="list-style-type: none"> <li>CEEDUMPs in controller (region) or servant (region). These dumps indicate that the runtime had an error</li> <li>JRAS error log entries, which can narrow the error down the exception to a specific function within the runtime</li> </ul>                                                                                                  | <ul style="list-style-type: none"> <li>CEEDUMPs</li> <li>JRAS error log entries and traces</li> </ul>                                                                                                                                                                                       |
| <b>Actions:</b>   | <ul style="list-style-type: none"> <li>Look at the minor code that is listed.</li> <li>Look for fixes that address similar symptoms or minor codes.</li> <li>System exceptions usually represent the detection of an unexpected error, and therefore (unless directed by the documentation of the minor code) will often require IBM assistance to identify the problem.</li> </ul> | <ul style="list-style-type: none"> <li>Look at your application for any sign of error.</li> <li>Look for system failures, such as a system exception in the controller (region). If you find a system exception, follow the steps to the left for diagnosing a system exception.</li> </ul> |

---

## Debugging applications that hang

|                        |                                                                                 |                                                                             |
|------------------------|---------------------------------------------------------------------------------|-----------------------------------------------------------------------------|
| <b>Possible cause:</b> | The WebSphere variable transaction_defaultTimeout might have a value too large. |                                                                             |
| <b>Caused by:</b>      | Loop in the application                                                         | JVM runs out of heap storage, if you are running Java in a servant (region) |

|                         |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |                                                                                                                               |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|
| <p><b>Look for:</b></p> | <ul style="list-style-type: none"> <li>• Environment variable that handles how long the application runs before timeout</li> <li>• Timeout-related minor codes: <ul style="list-style-type: none"> <li>– C9C21047</li> <li>– C9C2110F</li> <li>– C9C21110</li> <li>– C9C21111</li> <li>– C9C21112</li> <li>– C9C21113</li> <li>– C9C21114</li> <li>– C9C21190</li> <li>– C9C21191</li> <li>– C9C21192</li> <li>– C9C21809</li> <li>– C9C21892</li> <li>– C9C21893</li> <li>– C9C22013</li> </ul> </li> <li>• ABEND EC3, reason codes 0413002 through 04130007</li> <li>• resource messages on the console</li> </ul> <p><b>Example:</b> DB2 deadlock messages on the z/OS console</p> <ul style="list-style-type: none"> <li>• A wait beyond the timeout value length with no timeout</li> </ul> | <ul style="list-style-type: none"> <li>• Any error messages from JVM in the job log of the failed servant (region)</li> </ul> |
|-------------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------|

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>Actions:</b> | <ul style="list-style-type: none"> <li>• Analyze with IPCS to determine whether or not the servant (region) was looping (application code loop) or waiting (maybe the runtime failed). <ul style="list-style-type: none"> <li>– Use the DUMP command to get a console dump of the servant and its controller.</li> </ul> </li> <li>• If you were utilizing JRAS, look at the JRAS TRACE entries: <ul style="list-style-type: none"> <li>– If the application code was looping, you may see the same entry repeating.</li> </ul> </li> <li>• Ensure that TRACE writer is on and take a SVC dump at the approximate time of hang.</li> <li>• Use the display command to determine the state of the server.</li> </ul> | <ul style="list-style-type: none"> <li>• Through the Administrative console, set the WebSphere variable to debug the JVM; this setting passes information to the JVM and turns on the high-level messages for you to examine.</li> <li>• Look for error message or Java stack traces that might indicate an OUT_OF_MEMORY condition.</li> <li>• Use application monitoring tools, such as WebSphere Studio Application Monitor (WSAM) or Jinsight, to look for application memory leaks.</li> </ul> |
|-----------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|

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

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

---

## Debugging problems related to Java Message Service (JMS) support

You might encounter JMS-related errors in the WebSphere Application Server for z/OS environment. To debug these errors, use any of the various WebSphere console settings that controls the type of trace data collected.

**com.ibm.ejs.\*=all=enabled**

Turns on all container tracing

**com.ibm.ejs.j2c.\*=all=enabled**

Turns on all tracing for connector support in WebSphere Application Server for z/OS

**Messaging=all=enabled**

Collects trace data for the JMS and Message-driven bean (MDB) components of WebSphere Application Server for z/OS

If your installation configures WebSphere MQ to provide Java Message Service (JMS) support, you might need to use specific MQ tools for diagnosis:

- WebSphere variable *JMSApi=all=enabled*, which turns on all tracing for applications that use the JMS application programming interface (API).

- MQSC commands, which allow you to display information and perform other operations.
- The CSQUTIL utility program, which helps you to perform backup, restoration, and reorganization tasks, and to issue WebSphere MQ commands.

For more information about these diagnostic tools, refer to WebSphere MQ books, which are available through the Web site:

<http://www.ibm.com/software/ts/mqseries/library/>

The most useful for diagnostic information are:

- WebSphere MQ for z/OS Problem Determination, GC34-6054
- WebSphere MQ for z/OS Messages and Codes, SC34-6056
- WebSphere MQ Script (MQSC) Command Reference, SC34-6055
- WebSphere MQ for z/OS System Administration Guide, SC34-6053

All of these publications can be accessed from the IBM Publications Center.



---

## Chapter 10. Troubleshooting by component: What is not working?

This section provides troubleshooting information based on the task you were trying to accomplish when the problem occurred. To find more information about your problem, select a task category from the list below.

---

### Installation component troubleshooting tips

If you are having problems installing the WebSphere Application Server, follow these steps to resolve the problem:

- Browse the relevant log files for clues:
  - The main installation log file: *install\_dir/log.txt*.
  - IBM Http Server log: *install\_dir/ihs\_log.txt*.
  - The log file produced when the default application .ear file is installed is: *install\_dir/logs/installDefaultApplication.log*.
- Ensure that you have installed the correct level of dependent software, such as operating system version and revision level, by reviewing <http://www.ibm.com/software/webservers/appserv/doc/latest/prereq.html>.

If none of these steps solves the problem, check to see if the problem is identified and documented using the links in Diagnosing and fixing problems: Resources for learning.

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

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

---

### Administration and administrative console troubleshooting tips

In WebSphere Application Server products, administrative functions are supported by:

- The application server (such as server1) process in the base product
- The deployment manager (dmgr) process in the Network Deployment product

The process must be running to use the administrative console. The **wsadmin** command line utility has a local mode that you can use to perform administrative functions, even when the server process is not running.

When starting or stopping a server using a wsadmin interactive scripting session, you receive an exception indicating read timed out, for example:

```
WASX7015E: Exception running command: "$AdminControl startServer server1 Node1";
exception information: com.ibm.websphere.management.exception.ConnectorException
org.apache.soap.SOAPException: [SOAPException: faultCode=SOAP-ENV:Client; msg=Read
timed out; targetException=java.net.SocketTimeoutException: Read timed out]
```

This exception occurs because the timeout value is too small. To fix this, increase the timeout value specified by the `com.ibm.SOAP.requestTimeout` property in the `soap.client.props` file to 5 minutes.



If you have problems starting or using the administrative console or wsadmin utility, verify that the supporting server process is started and that it is healthy.

- For the base product, look at these files:
  - *install\_root/logs/server/startServer.log* for the message that indicates that the server started successfully: **ADMU3000I: Server server1 open for e-business; process id is nnnn..**
  - the server log files for the message that indicates that the server started successfully: **WSVR0001I: Server server open for e-business.**
- For the Network Deployment product, look at these files:
  - *install\_root/logs/dmgr/startServer.log* for the message that indicates that the server started successfully: **ADMU3000I: Server dmgr open for e-business; process id is nnnn..**
  - the server log files for this message indicating that the server started successfully: **WSVR0001I: Server dmgr open for e-business.**
- For the z/OS product, check the job output.
- Look up any error messages in these files in the message reference table. Select the **Reference** view in the information center navigation, then click **Messages**.
- A message like **WASX7213I: This scripting client is not connected to a server process** when trying to start wsadmin indicates that either the server process is not running, the host machine where it is running is not accessible, or that the port or server name used by wsadmin is incorrect.
- Verify that you are using the right port number to communicate with the administrative console or wsadmin server using the following steps:
  - Look in the the server log files.
  - The line **ADMC0013I: SOAP connector available at port nnnn** indicates the port that the server is using to listen for wsadmin functions.
  - The property **com.ibm.ws.scripting.port** in the *install\_root/properties/wsadmin.properties* file controls the port used by wsadmin to send requests to the server. If it is different from the value shown in the the server log files, either change the port number in the *wsadmin.properties* file, or specify the correct port number when starting wsadmin by using the **-port port\_number** property on the command line.
  - The message **SRVE0171I: Transport http is listening on port nnnn (default 9090)** indicates the port the server uses to listen for administrative console requests. If it is different than the one specified in the URL for the administrative console, change the URL in the browser to the correct value. The default value is `http://localhost:9090/admin`.
- Use the TCP/IP **ping** command to test that the hostname where the application server or deployment manager is executing, is reachable from the system where the browser or wsadmin program are being used. If you are able to ping the hostname, this indicates that there are no firewall or connectivity issues.
- If the host where the application server or deployment manager is running is remote to the machine from which the client browser or wsadmin command is running, ensure that the appropriate hostname parameter is correct:
  - The hostname in the browser URL for the console.
  - The **-host hostname** option of the wsadmin command that is used to direct wsadmin to the right server

If none of these steps solves the problem, see if the specific problem you are having is addressed in the Installation completes but the administrative console does not start topic. Check to see if the problem has been identified and documented using the links in the Diagnosing and fixing problems: Resources for learning topic.

For current information available from IBM Support on known problems and their resolution, see the following topics on the IBM support page:

- Administrative interface
- System management

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

- Administrative interface
- System management

---

## Web Container troubleshooting tips

If you are having problems starting a Web module, or accessing resources within a particular Web module:

- For specific problems that can cause servlets, HTML files, and JavaServer Pages (JSP) files not to be served, see [Web resource \(JSP file, servlet, HTML file, image\) does not display](#).

If none of these steps fixes your problem, check to see if the problem has been identified and documented by looking at the available online support (hints and tips, technotes, and fixes).

For current information available from IBM Support on known problems and their resolution, see the [IBM Support page](#).

IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the [IBM Support page](#).

---

## JDBC and data source troubleshooting tips

To see whether your specific problem has been addressed, review the [Cannot access a data source](#) topic.

For current information available from IBM Support on known problems and their resolution, see the [IBM Support page](#).

IBM Support has documents that can save you time gathering information needed to resolve this problem. Before opening a PMR, see the [IBM Support page](#).

---

## HTTP plug-in component troubleshooting tips

If you are having problems with the HTTP plug-in component - the component which sends requests from your HTTP server, such as IBM HTTP Server, Apache, Domino, iPlanet, or IIS, to the WebSphere Application Server, try these steps:

- Review the file `install_dir/logs/http_plugin.log` for clues. Look up any error or warning messages in the message table.
- Review your HTTP server's error and access logs to see if the HTTP server is having a problem:
  - IBM HTTP Server and Apache: `access.log` and `error.log`.
  - Domino web server: `httpd-log` and `httpd-error`.
  - iPlanet: `access` and `error`.
  - IIS: `timedatestamp.log`.

If these files don't reveal the cause of the problem, follow these additional steps.

## Plug-in Problem Determination Steps

The plug-in provides very readable tracing which can be beneficial in helping to figure out the problem. By setting the **LogLevel** attribute in the `config/plugin-cfg.xml` file to **Trace**, you can follow the request processing to see what is going wrong. At a high level:

1. The plug-in gets a request.
2. The plug-in checks the routes defined in the `plugin-cfg.xml` file.
3. It finds the server group.
4. It finds the server.
5. It picks the transport protocol, usually HTTP.
6. It sends the request.
7. It reads the response.
8. It writes it back to the client.

You can see this very clearly by reading through the trace for a single request:

- The first step is to determine if the plug-in has loaded into the HTTP server successfully.
  - Check to make sure `thehttp_plugin.log` has been created.
  - If it has, look in it to see if any error messages indicate some sort of failure that took place during plug-in initialization. If no errors are found look for the following stanza, which indicates that the plug-in started normally. Ensure that the timestamps for the messages correspond to the time you started the Web server:

```
[Thu Jul 11 10:59:15 2002] 0000009e 000000b1 - PLUGIN: ---System Information---
[Thu Jul 11 10:59:15 2002] 0000009e 000000b1 - PLUGIN: Bld date: Jul 3 2002, 15:35:09
[Thu Jul 11 10:59:15 2002] 0000009e 000000b1 - PLUGIN: Web server: IIS
[Thu Jul 11 10:59:15 2002] 0000009e 000000b1 - PLUGIN: Hostname = SWEETTJ05
[Thu Jul 11 10:59:15 2002] 0000009e 000000b1 - PLUGIN: OS version 4.0, build 1381,
 'Service Pack 6'
[Thu Jul 11 10:59:15 2002] 0000009e 000000b1 - PLUGIN: -----
```

- Some common errors are:

### **lib\_security: loadSecurityLibrary: Failed to load gsk library**

The GSKit did not get installed or the installation is corrupt. If the GSKit did not get installed you can determine this by searching for the file `gsk5ssl.dll` on all drives for Win32 or see if there are any `libgsk5*.so` files in `/usr/lib` on UNIX. Try reinstalling the plug-in to see if you can get the GSKit to install in order to fix this.

### **ws\_transport: transportInitializeSecurity: Keyring wasn't set**

The HTTPS transport defined in the configuration file was prematurely terminated and did not contain the Property definitions for the keyring and stashfile. Check your XML syntax for the line number given in the error messages that follow this one to make sure the Transport element contains definitions for the keyring and stashfiles before it is terminated.

- If `thehttp_plugin.log` is not created, check the Web server error log to see if any plug-in related error messages have been logged there that indicate why the plug-in is failing to load. Typical causes of this can include failing to correctly configure the plug-in with the Web server environment. Check the documentation for configuring the Web server you are trying to use with the Web server plug-in.
- Determine whether there are network connection problems with the plug-in and the various application servers defined in the configuration. Typically you will see the following message when this is the case:

```
ws_common: webspHEREGetStream: Failed to connect to app server, OS
err=%d
```

Where %d is an OS specific error code related to why the connect() call failed. This can happen for a variety of reasons.

- Ping the machines to make sure they are properly connected to the network. If the machines can't be pinged then there is no way the plug-in will be able to contact them. Possible reasons for this include:
  - Firewall policies limiting the traffic from the plug-in to the app server.
  - The machines are not on the same network.
- If you are able to ping the machines then the likely cause of the problem is that the port is not active. This could be because the application server or cluster has not been started or the application server has gone down for some reason. You can test this by hand by trying to telnet into the port that the connect is failing on. If you cannot telnet into the port the application server is not up and that problem needs to be resolved before the plug-in will be able to connect successfully.
- Determine whether other activity on the machines where the servers are installed is impairing the server's ability to service a request. Check the processor utilization as measured by the task manager, processor ID, or some other outside tool to see if it:
  - Is not what was expected.
  - Is erratic rather than a constant.
  - Shows that a newly added member of the cluster is not being utilized.
  - Shows that a failing member that has been fixed is not being utilized.
- Check the administrative console to ensure that the application servers are started. View the administrative console for error messages.
- In the administrative console, select the problem application server and view its installed applications to verify that they are started.

If none of these steps solves the problem:

- For specific problems that can cause web pages and their contents not to display, see Web resource (JSP file, servlet, html file, image, etc) will not display in the information center.
- Check to see if the problem has been identified and documented using the links in Diagnosing and fixing problems: Resources for learning.

For current information available from IBM Support on known problems and their resolution, see the following topics on the IBM support page:

- HTTP transport
- HTTP plugin
- HTTP plugin remote install

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

- HTTP transport
- HTTP plugin
- HTTP plugin remote install

---

## HTTP session manager troubleshooting tips

If you are having problems creating or using HTTP sessions with your Web application hosted by WebSphere Application Server, here are some steps to take:

- View the logs for the application server which hosts the problem application:
  - first, look at messages written while each application is starting. They will be written between the following two messages:

Starting application: *application*

.....

Application started: *application*

- Within this block, look for any errors or exceptions containing a package name of `com.ibm.ws.webcontainer.httpsession`. If none are found, this is an indication that the session manager started successfully.
- Error "**SRVE0054E: An error occurred while loading session context and Web application**" indicates that SessionManager didn't start properly for a given application.
- Look within the logs for any Session Manager related messages. These messages will be in the format `SESNxxxxE` and `SESNxxxxW` for errors and warnings, respectively, where `xxxx` is a number identifying the precise error. Look up the extended error definitions in the Session Manager message table.
- See Best practices for using HTTP Sessions.
- Alternatively, a special servlet can be invoked that displays the current configuration and statistics related to session tracking.
  - Servlet name: **`com.ibm.ws.webcontainer.httpsession.IBMTrackerDebug`**.
  - It can be invoked from any web module which is enabled to serve by class name. For example, using `default_app`,  
**`http://localhost:9080/servlet/com.ibm.ws.webcontainer.httpsession.IBMTrackerDebug`**.
  - If you are viewing the module via the `serve-by-class-name` feature, be aware that it may be viewable by anyone who can view the application. You may wish to map a specific, secured URL to the servlet instead and disable the `serve-servlets-by-classname` feature.
- If you are using **database-based persistent sessions**, look for problems related to the **data source** the Session Manager relies on to keep session state information. For details on diagnosing database related problems see Errors accessing a `datasource` or `connection pool`

### Error message SRVE0079E Servlet host not found after you define a port

Error message SRVE0079E can occur after you define the port in WebContainer > HTTP Transports for a server, indicating that you do not have the port defined in your virtual host definitions. To define the port,

1. On the administrative console, go to Environment > Virtual Hosts > `default_host`> Host Aliases> New
2. Define the new port on host "\*"

### The application server gets EC3 - 04130007 ABENDs

To prevent an EC3 - 04130007 abend from occurring on the application server, change the HTTP Output timeout value. The custom property `ConnectionResponseTimeout` specifies the maximum number of seconds the HTTP port for an individual server can wait when trying to read or write data. For instructions on how to set `ConnectionResponseTimeout`, see HTTP transport custom properties.

If none of these steps fixes your problem, check to see if the problem has been identified and documented by looking at the available online support (hints and tips, technotes, and fixes).

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

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

---

## Naming services component troubleshooting tips

"Naming" is a J2EE service which publishes and provides access to resources such as connection pools, enterprise beans, message listeners, etc, to client processes. If you have problems in accessing a resource which otherwise appears to be healthy, the naming service might be involved. To investigate problems with the WebSphere Application Server Naming service:

- With WebSphere Application Server running, run the **dumpNameSpace** command for Windows systems, or the **dumpNameSpace.sh** command for UNIX and z/OS systems, and pipe, redirect, or "more" the output so that it is easily viewed. This command results in a display of all objects in the WebSphere Application Server namespace, including the directory path and object name.
- If the object a client needs to access does not appear, use the administrative console to verify that:
  - The server hosting the target resource is started.
  - The Web module or EJB container, if applicable, hosting the target resource is running.
  - The JNDI name of the target resource is correct and updated.
  - If the problem resource is remote, that is, not on the same node as the Name Server node, that the JNDI name is fully qualified, including the host name. This is especially applicable to Network Deployment configurations
- If you see an exception that appears to be CORBA related ("CORBA" appears as part of the exception name) look for a naming-services-specific CORBA minor code, further down in the exception stack, for information on the real cause of the problem. For a list of naming service exceptions and explanations, see the class `com.ibm.websphere.naming.WsnCorbaMinorCodes` in the Javadoc topic in the information center.

If none of these steps solve the problem:

- For specific problems that can cause access to named object hosted in WebSphere Application Server to fail, see Cannot look up an object hosted by WebSphere Application Server from a servlet, JSP file, or other client.
- Check to see if the problem has been identified and documented using the links in Diagnosing and fixing problems: Resources for learning.

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

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

---

## Messaging component troubleshooting tips

### Problems deploying or executing applications which use the WebSphere Application Server messaging capabilities

If you are having problems deploying or executing applications which use the WebSphere Application Server messaging capabilities, review these articles in the WebSphere Application Server information center:

- Troubleshooting WebSphere Messaging
- Troubleshooting message-driven beans
- Troubleshooting transactions



### **On a Linux platform, embedded messaging does not get removed after uninstalling the product**

On a Linux platform, you may find that embedded messaging is not removed when you uninstall the product silently using the `SetRetainMQToTrue.active=false` option.

To prevent this problem, either issue the `uninstall` command from a local machine, or use `ssh` instead of `telnet` to logon to the machine where you issue the `uninstall` command.

To correct this problem after logging in, issue the following command for a real root login: `su -`

If none of these steps solves the problem, check to see if the problem has been identified and documented using the links in *Diagnosing and fixing problems: Resources for learning*.

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

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

---

## **Universal Discovery, Description, and Integration, Web Service, and SOAP component troubleshooting tips**

If you are having problems deploying or executing applications that use WebSphere Application Server Web Services, Universal Discovery, Description, and Integration (UDDI), or SOAP, try these steps:

- Review the troubleshooting documentation for messaging in the information center:
  - WSIF troubleshooting tips
- Investigate the following areas for SOAP-related problems:
  - View the error log of the HTTP server to which the SOAP request is sent.
  - Browse the Web site <http://xml.apache.org/soap/> for FAQs and known SOAP issues.

If none of these steps solves the problem, check to see if the problem has been identified and documented using the links in *Diagnosing and fixing problems: Resources for learning*.

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

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

---

## **Enterprise bean and EJB container troubleshooting tips**

If you are having problems starting an EJB container, or encounter error messages or exceptions that appear to be generated on by an EJB container, follow these steps to resolve the problem:

- Browse the relevant log files for clues:



- Use the Administrative Console to verify that the application server which hosts the container is running.
- Browse the logs for the application server which hosts the container. Look for the message **server *server\_name* open for e-business** in the server log files. If it does not appear, or if you see the message **problems occurred during startup**, browse the server log files for details.

If none of these steps solves the problem, check to see if the problem is identified and documented using the links in Diagnosing and fixing problems: Resources for learning.

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

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

---

## Cannot restart the Deployment Manager monitoring policy

The Deployment Manager monitoring policy is permanently set to STOPPED and cannot be changed. Therefore, if the Deployment Manager fails it will never be restarted by WebSphere Application Server monitoring.

To correct the problem, use the z/OS Automatic Restart Facility (ARM) to monitor and restart the Deployment Manager. See z/OS Automatic Restart Facility for information on ARM.

---

## Security components troubleshooting tips

This document explains basic resources and steps for diagnosing security related issues in the WebSphere Application Server, including:

- What log files to look at and what to look for in them.
- A general approach to isolating and resolving security problems.
- When and how to enable security-related trace.

The following security-related problems are addressed elsewhere in the information center:

- Errors and access problems after enabling security
- Errors after enabling SSL, or SSL-related error messages

If none of these steps solves the problem, check to see if the problem has been identified and documented using the links in Diagnosing and fixing problems: Resources for learning.

**Note:** for an overview of WebSphere Application Server security components such as z/SAS and how they work, see Getting started with security.

When troubleshooting the security component, browse the SDSF logs for the server that hosts the resource you are trying to access. The following is a sample of messages you would expect to see from a server in which the security service has started successfully:

Messages beginning with BBOM0001I are related to zOS specific implementations of z/SAS and CSIV2. They appear in both the controller and servant but are only applicable in the controller.

```
BBOM0001I com_ibm_Server_Security_Enabled: 1.
```

```

BBOM0001I com_ibm_CSI_claimTLClientAuthenticationSupported: 1.
BBOM0001I com_ibm_CSI_claimTLClientAuthenticationRequired: 0.
BBOM0001I com_ibm_CSI_claimTransportAssocSSLTLSSupported: 1.
BBOM0001I com_ibm_CSI_claimTransportAssocSSLTLSRequired: 0.
BBOM0001I com_ibm_CSI_claimMessageConfidentialityRequired: 0.
BBOM0001I com_ibm_CSI_claimClientAuthenticationSupported: 1.
BBOM0001I com_ibm_CSI_claimClientAuthenticationRequired: 0.
BBOM0001I com_ibm_CSI_claimClientAuthenticationtype:
SAFUSERIDPASSWORD.
BBOM0001I com_ibm_CSI_claimIdentityAssertionTypeSAF: 0.
BBOM0001I com_ibm_CSI_claimIdentityAssertionTypeDN: 0.
BBOM0001I com_ibm_CSI_claimIdentityAssertionTypeCert: 0.
BBOM0001I com_ibm_CSI_claimMessageIntegritySupported: NOT SET,DEFAULT=1.
BBOM0001I com_ibm_CSI_claimMessageIntegrityRequired: NOT SET,DEFAULT=1.
BBOM0001I com_ibm_CSI_claimStateful: 1.
BBOM0001I com_ibm_CSI_claimSecurityLevel: HIGH.
BBOM0001I com_ibm_CSI_claimSecurityCipherSuiteList: NOT SET.
BBOM0001I com_ibm_CSI_claimKeyringName: WASKeyring.
BBOM0001I com_ibm_CSI_claim_ssl_sys_v2_timeout: NOT SET, DEFAULT=100.
BBOM0001I com_ibm_CSI_claim_ssl_sys_v3_timeout: 600.
BBOM0001I com_ibm_CSI_performTransportAssocSSLTLSSupported: 1.
BBOM0001I security_sslClientCerts_allowed: 0.
BBOM0001I security_kerberos_allowed: 0.
BBOM0001I security_userid_password_allowed: 0.
BBOM0001I security_userid_passticket_allowed: 1.
BBOM0001I security_assertedID_IBM_accepted: 0.
BBOM0001I security_assertedID_IBM_sent: 0.
BBOM0001I nonauthenticated_clients_allowed: 1.
BBOM0001I security_remote_identity: WSGUEST.
BBOM0001I security_local_identity: WSGUEST.
BBOM0001I security_EnableRunAsIdentity: 0.

```

Messages beginning with BB000222I are common to java WebSphere security. They appear in both the controller and servant but are applicable to the servant.

```

BB000222I SECJ0240I: Security service initialization completed successfully
BB000222I SECJ0215I: Successfully set JAAS login provider
configuration class to com.ibm.ws.security.auth.login.Configuration.
BB000222I SECJ0136I: Custom
Registry:com.ibm.ws.security.registry.zOS.SAFRegistryImpl has been initialized
BB000222I SECJ0157I: Loaded Vendor AuthorizationTable:
com.ibm.ws.security.core.SAFAuthorizationTableImpl
BB000222I SECJ0243I: Security service started successfully
BB000222I SECJ0210I: Security enabled true

```

## General approach for troubleshooting security-related issues

When troubleshooting security-related problems, the following questions are very helpful and should be considered:

### Does the problem occur when security is disabled?

This is a good litmus test to determine that a problem is security related. However, just because a problem only occurs when security is enabled does not always make it a security problem. More troubleshooting is necessary to ensure the problem is really security-related.

### Did security appear to initialize properly?

A lot of security code is visited during initialization. So you will likely see problems there first if the problem is configuration related. The following sequence of messages generated in the SDSF active log indicate normal code initialization of an application server. Non-security messages have been removed from the sequence that follows. This sequence will vary based on the configuration, but the messages are similar:

Trace: 2003/08/25 13:06:31.034 01 t=9EA930 c=UNK key=P8 (13007002)  
 FunctionName: com.ibm.ws.security.auth.login.Configuration  
 SourceId: com.ibm.ws.security.auth.login.Configuration  
 Category: AUDIT  
 ExtendedMessage: SECJ0215I: Successfully set JAAS login provider configuration class to com.ibm.ws.security.auth.login.Configuration.

Trace: 2003/08/25 13:06:31.085 01 t=9EA930 c=UNK key=P8 (13007002)  
 FunctionName: com.ibm.ws.security.core.SecurityDM  
 SourceId: com.ibm.ws.security.core.SecurityDM  
 Category: INFO  
 ExtendedMessage: SECJ0231I: The Security component's FFDC Diagnostic Module com.ibm.ws.security.core.SecurityDM registered success  
 fully: true.

Trace: 2003/08/25 13:06:31.086 01 t=9EA930 c=UNK key=P8 (0000000A)  
 Description: Log Boss/390 Error  
 from filename: ./bborjtr.cpp  
 at line: 812  
 error message: BB000222I SECJ0231I: The Security component's FFDC Diagnostic Module com.ibm.ws.security.core.SecurityDM registered successfully: true.

Trace: 2003/08/25 13:06:32.426 01 t=9EA930 c=UNK key=P8 (13007002)  
 FunctionName: com.ibm.ws.security.core.SecurityComponentImpl  
 SourceId: com.ibm.ws.security.core.SecurityComponentImpl  
 Category: INFO  
 ExtendedMessage: SECJ0309I: Java 2 Security is disabled.

Trace: 2003/08/25 13:06:32.427 01 t=9EA930 c=UNK key=P8 (0000000A)  
 Description: Log Boss/390 Error  
 from filename: ./bborjtr.cpp  
 at line: 812  
 error message: BB000222I SECJ0309I: Java 2 Security is disabled.

Trace: 2003/08/25 13:06:32.445 01 t=9EA930 c=UNK key=P8 (13007002)  
 FunctionName: com.ibm.ws.security.core.SecurityComponentImpl  
 SourceId: com.ibm.ws.security.core.SecurityComponentImpl  
 Category: INFO  
 ExtendedMessage: SECJ0212I: WCCM JAAS configuration information successfully pushed to login provider class.

Trace: 2003/08/25 13:06:32.445 01 t=9EA930 c=UNK key=P8 (0000000A)  
 Description: Log Boss/390 Error  
 from filename: ./bborjtr.cpp  
 at line: 812  
 error message: BB000222I SECJ0212I: WCCM JAAS configuration information successfully pushed to login provider class.

Trace: 2003/08/25 13:06:32.459 01 t=9EA930 c=UNK key=P8 (13007002)  
 FunctionName: SecurityComponentImpl  
 SourceId: SecurityComponentImpl  
 Category: WARNING  
 ExtendedMessage: BBOS1000W LTPA or ISCF are configured as the authentication mechanism but SSO is disabled.

Trace: 2003/08/25 13:06:32.459 01 t=9EA930 c=UNK key=P8 (0000000A)  
 Description: Log Boss/390 Error  
 from filename: ./bborjtr.cpp  
 at line: 824  
 error message: BBOS1000W LTPA or ISCF are configured as the authentication mechanism but SSO is disabled.

Trace: 2003/08/25 13:06:32.463 01 t=9EA930 c=UNK key=P8 (13007002)  
 FunctionName: com.ibm.ws.security.core.SecurityComponentImpl  
 SourceId: com.ibm.ws.security.core.SecurityComponentImpl  
 Category: INFO  
 ExtendedMessage: SECJ0240I: Security service initialization completed successfully

Trace: 2003/08/25 13:06:32.463 01 t=9EA930 c=UNK key=P8 (0000000A)  
 Description: Log Boss/390 Error  
 from filename: ./bborjtr.cpp  
 at line: 812  
 error message: BB000222I SECJ0240I: Security service initialization completed successfully

```

Trace: 2003/08/25 13:06:39.718 01 t=9EA930 c=UNK key=P8 (13007002)
 FunctionName: com.ibm.ws.security.registry.UserRegistryImpl
 SourceId: com.ibm.ws.security.registry.UserRegistryImpl
 Category: AUDIT
 ExtendedMessage: SECJ0136I: Custom Registry:
com.ibm.ws.security.registry.zOS.SAFRegistryImpl has been initialized
Trace: 2003/08/25 13:06:41.967 01 t=9EA930 c=UNK key=P8 (13007002)
 FunctionName: com.ibm.ws.security.core.WSAccessManager
 SourceId: com.ibm.ws.security.core.WSAccessManager
 Category: AUDIT
 ExtendedMessage: SECJ0157I: Loaded Vendor AuthorizationTable:
com.ibm.ws.security.core.SAFAuthorizationTableImpl
Trace: 2003/08/25 13:06:43.136 01 t=9EA930 c=UNK key=P8 (13007002)
 FunctionName: com.ibm.ws.security.role.RoleBasedAuthorizerImpl
 SourceId: com.ibm.ws.security.role.RoleBasedAuthorizerImpl
 Category: AUDIT
 ExtendedMessage: SECJ0157I: Loaded Vendor AuthorizationTable:
com.ibm.ws.security.core.SAFAuthorizationTableImpl
Trace: 2003/08/25 13:06:43.789 01 t=9EA930 c=UNK key=P8 (13007002)
 FunctionName: com.ibm.ws.security.core.SecurityComponentImpl
 SourceId: com.ibm.ws.security.core.SecurityComponentImpl
 Category: INFO
 ExtendedMessage: SECJ0243I: Security service started successfully
Trace: 2003/08/25 13:06:43.789 01 t=9EA930 c=UNK key=P8 (0000000A)
 Description: Log Boss/390 Error
 from filename: ./bborjtr.cpp
 at line: 812
 error message: BB000222I SECJ0243I: Security service started successfully
Trace: 2003/08/25 13:06:43.794 01 t=9EA930 c=UNK key=P8 (13007002)
 FunctionName: com.ibm.ws.security.core.SecurityComponentImpl
 SourceId: com.ibm.ws.security.core.SecurityComponentImpl
 Category: INFO
 ExtendedMessage: SECJ0210I: Security enabled true
Trace: 2003/08/25 13:06:43.794 01 t=9EA930 c=UNK key=P8 (0000000A)
 Description: Log Boss/390 Error
 from filename: ./bborjtr.cpp
 at line: 812
 error message: BB000222I SECJ0210I: Security enabled true
Trace: 2003/08/25 13:07:06.474 01 t=9EA930 c=UNK key=P8 (13007002)
 FunctionName: com.ibm.ws.security.core.WSAccessManager
 SourceId: com.ibm.ws.security.core.WSAccessManager
 Category: AUDIT
 ExtendedMessage: SECJ0157I: Loaded Vendor AuthorizationTable:
com.ibm.ws.security.core.SAFAuthorizationTableImpl
Trace: 2003/08/25 13:07:09.315 01 t=9EA930 c=UNK key=P8 (13007002)
 FunctionName: com.ibm.ws.security.core.WSAccessManager
 SourceId: com.ibm.ws.security.core.WSAccessManager
 Category: AUDIT
 ExtendedMessage: SECJ0157I: Loaded Vendor AuthorizationTable:
com.ibm.ws.security.core.SAFAuthorizationTableImpl

```

### **Is this a distributed security problem or a local security problem?**

- If the problem is local, that is the code involved does not make a remote method invocation, then troubleshooting is isolated to a single process. It is important to know when a problem is local versus distributed since the behavior of the ORB, among other components, is different between the two. Once a remote method invocation takes place, an entirely different security code path is entered.
- When you know that the problem involves two or more servers, the techniques of troubleshooting change. You will need to trace all servers involved simultaneously so that the trace shows the client and server sides of the problem. Try to make sure the timestamps on all machines match as closely as possible so that you can find the request and reply pair from two different processes. .

### **Is the problem related to authentication or authorization?**

Most security problems fall under one of these two categories.

Authentication is the process of determining who the caller is. Authorization is the process of validating that the caller has the proper authority to invoke the requested method. When authentication fails, typically this is related to either the authentication protocol, authentication mechanism or user registry. When authorization fails, this is usually related to the application bindings from assembly and/or deployment and to the caller's identity who is accessing the method and the roles required by the method.

### **Is this a Web or EJB request?**

Web requests have a completely different code path than EJB requests. Also, there are different security features for Web requests than for EJB requests, requiring a completely different body of knowledge to resolve. For example, when using the LTPA authentication mechanism, the Single SignOn feature is available for Web requests but not for EJB requests. Web requests involve HTTP header information not required by EJB requests due to the protocol differences. Also, the Web container (or servlet engine) is involved in the entire process. Any of these components could be involved in the problem and all should be considered during troubleshooting, based on the type of request and where the failure occurs.

Secure EJB requests are passed from the controller to the servant. Web requests are mostly ignored by the controller. As a result, EJB requests are first processed and authenticated by the z/SAS or CSIv2 layers of security. Authorization is done by the servant. If an authentication failure occurs, the z/SAS type level of tracing must be turned on to diagnose the problem. Other problems can be diagnosed using the WebSphere Application Server component tracing (CTRACE) facility.

### **Does the problem seem to be related to the Secure Sockets Layer (SSL)?**

The Secure Socket Layer (SSL) is a totally distinct separate layer of security. Troubleshooting SSL problems are usually separate from troubleshooting authentication and/or authorization problems. There are many things to consider. Usually, SSL problems are first time setup problems because the configuration can be difficult. Each client must contain the server's signer certificate. During mutual authentication, each server must contain the client's signer certificate. Also, there can be protocol differences (SSLv3 vs. TLS), and listener port problems related to stale IORs (i.e., IORs from a server reflecting the port prior to the server restarting).

In z/OS, two variations of SSL are used. To determine the cause of an SSL problem on z/OS, you will have to be aware of what protocol is being used. System SSL is used by the IIOP and HTTPS protocols. Java Secure Socket Extension (JSSE) is used by all other protocols, for example, Simple Object Access Protocol (SOAP). System SSL requests are handled in the controller and are used by z/SAS and CSIv1 security. JSSE is predominately used by the servant, but there are cases where it is used in the controller as well.

For SSL problems, we sometimes request an SSL trace to determine what is happening with the SSL handshake. The SSL handshake is the process which occurs when a client opens a socket to a server. If anything goes wrong with the key exchange, cipher exchange, etc. the handshake will fail and thus the socket is invalid. Tracing JSSE (the SSL implementation used in WebSphere Application Server) involves the following steps:

- Set the following system property on the client and server processes:  
-Djavax.net.debug=true. For the server, add this to the Generic JVM Arguments property of the Java virtual machine settings page.
- Recreate the problem. The SDSF active log of both processes should contain the JSSE trace. You will find trace similar to the following:

```

JSSEContext: handleConnection[Socket
[addr=boss0106.plex1.l2.ibm.com/9.38.48.108,port=2139,localport=8878]]
JSSEContext: handleConnection[Socket
[addr=boss0106.plex1.l2.ibm.com/9.38.48.108,port=2140,localport=8878]]
TrustManagerFactoryImpl: trustStore is :
/WebSphere/V5R0M0/AppServer/etc/DummyServerTrustFile.jks
TrustManagerFactoryImpl: trustStore type is : JKS
TrustManagerFactoryImpl: init truststore
JSSEContext: handleConnection[Socket
[addr=boss0106.plex1.l2.ibm.com/9.38.48.108,port=2142,localport=8878]]
KeyManagerFactoryImpl: keyStore is :
/WebSphere/V5R0M0/AppServer/etc/DummyServerKeyFile.jks
KeyManagerFactoryImpl: keyStore type is : JKS
KeyManagerFactoryImpl: init keystore
KeyManagerFactoryImpl: init keystore
JSSEContext: handleConnection[Socket
[addr=boss0106.plex1.l2.ibm.com/9.38.48.108,port=2143,localport=8878]]
JSSEContext: handleSession[Socket
[addr=BOSSXXXX.PLEX1.L2.IBM.COM/9.38.48.108,port=8879,localport=2145]]
JSSEContext: confirmPeerCertificate
[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM/9.38.48.108,port=8879,
localport=2145]]
X509TrustManagerImpl: checkServerTrusted
X509TrustManagerImpl: Certificate [
[
Version: V3
Subject: CN=jserver, OU=SWG, O=IBM, C=US
Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4
0 Key: IBMJCE RSA Public Key:
modulus:
10094996692239509074796828756118539107568369566313889955538950668
6622953008589748001058216362638201577071902071311277365773252660799
128781182947273802312699983556527878615792292244995317112436562491
489904381884265119355037731265408654007388863303101746314438337601
264540735679944205391693242921331551342247891
public exponent:
65537
0 Validity: [From: Fri Jun 21 20:08:18 GMT 2002,
To: Thu Mar 17 20:08:18 GMT 2005]
Issuer: CN=jserver, OU=SWG, O=IBM, C=US
SerialNumber: [3d1387b2]
0]
Algorithm: [MD5withRSA]
Signature:
0000: 54 DC B5 FA 64 C9 CD FE B3 EF 15 22 3D D0 20 31 T...d....."=. 1
0010: 99 F7 A7 86 F9 4C 82 9F 6E 4B 7B 47 18 2E C6 25L.nK.G...%
0020: 5B B2 9B 78 D8 76 5C 82 07 95 DD B8 44 62 02 62 [...x.v\.....Db.b
0030: 60 2A 0A 6D 4F B9 0A 98 14 27 E9 BB 1A 84 8A D1 ~*.m0....'.....
0040: C2 22 AF 70 9E A5 DF A2 FD 57 37 CE 3A 63 1B EB ."p.....W7.:c..
0050: E8 91 98 9D 7B 21 4A B5 2C 94 FC A9 30 C2 74 72!J,....0.tr
0060: 95 01 54 B1 29 E7 F8 9E 6D F3 B5 D7 B7 D2 9E 9B ..T.)...m.....
0070: 85 D8 E4 CF C2 D5 3B 64 F0 07 17 9E 1E B9 2F 79;d...../y
0]
X509TrustManagerImpl: Certificate [
[
Version: V3
Subject: CN=jserver, OU=SWG, O=IBM, C=US
Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4
0 Key: IBMJCE RSA Public Key:
modulus:
1009499669223950907479682875611853910756836956631388995553895066866

```



```

22953008589748001058216362638201577071902071311277365773252660799
1287811829472738023126999835565278786157922922449953171124365624914
89904381884265119355037731265408654007388863303101746314438337601
264540735679944205391693242921331551342247891
public exponent:
65537
0 Validity: [From: Fri Jun 21 20:08:18 GMT 2002,
 To: Thu Mar 17 20:08:18 GMT 2005]
 Issuer: CN=jserver, OU=SWG, O=IBM, C=US
 SerialNumber: [3d1387b2]
0]
 Algorithm: [MD5withRSA]
 Signature:
0000: 54 DC B5 FA 64 C9 CD FE B3 EF 15 22 3D D0 20 31 T...d....."=. 1
0010: 99 F7 A7 86 F9 4C 82 9F 6E 4B 7B 47 18 2E C6 25 L..nK.G...%
0020: 5B B2 9B 78 D8 76 5C 82 07 95 DD B8 44 62 02 62 [...x.v\.....Db.b
0030: 60 2A 0A 6D 4F B9 0A 98 14 27 E9 BB 1A 84 8A D1 `*.m0....'.....
0040: C2 22 AF 70 9E A5 DF A2 FD 57 37 CE 3A 63 1B EB ".p.....W7.:c..
0050: E8 91 98 9D 7B 21 4A B5 2C 94 FC A9 30 C2 74 72 !J,....0.tr
0060: 95 01 54 B1 29 E7 F8 9E 6D F3 B5 D7 B7 D2 9E 9B ..T.)...m.....
0070: 85 D8 E4 CF C2 D5 3B 64 F0 07 17 9E 1E B9 2F 79 ;d...../y
0]
 JSSEContext: handleConnection[Socket[addr=boss0106.plex1.l2.ibm.com
/9.38.48.108,port=2144,localport=8878]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2145]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2146]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2147]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2148]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2149]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2150]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2151]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2152]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2153]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2154]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2155]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2156]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2157]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2158]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2159]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2160]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2161]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2162]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2163]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2164]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2165]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM

```



```

/9.38.48.108,port=8879,localport=2166]]

JSSEContext: handleSession[Socket[addr=boss0106.plex1.12.ibm.com
/9.38.48.108,port=9443,localport=2167]]
JSSEContext: confirmPeerCertificate[Socket[addr=boss0106.plex1.12.ibm.com
/9.38.48.108,port=9443,localport=2167]]
X509TrustManagerImpl: checkServerTrusted
X509TrustManagerImpl: Certificate [
[
 Version: V3
 Subject: CN=WAS z/OS Deployment Manager, O=IBM
 Signature Algorithm: SHA1withRSA, OID = 1.2.840.113549.1.1.5
0 Key: IBMJCE RSA Public Key:
modulus:
12840948267119651469312486548020957441946413494498370439558603901582589
8755033448419534105183133064366466828741516428176579440511007
6258795528749232737808897160958348495006972731464152299032614592135114
19361539962555997136085140591098259345625853617389396340664766
649957749527841107121590352429348634287031501
public exponent:
65537
0 Validity: [From: Fri Jul 25 05:00:00 GMT 2003,
 To: Mon Jul 26 04:59:59 GMT 2004]
 Issuer: CN=WAS CertAuth, C=US
 SerialNumber: [02]
0Certificate Extensions: 3
[1]: ObjectId: 2.16.840.1.113730.1.13 Criticality=false
Extension unknown: DER encoded OCTET string =
0000: 04 3C 13 3A 47 65 6E 65 72 61 74 65 64 20 62 79 .<.:Generated by
0010: 20 74 68 65 20 53 65 63 75 72 65 57 61 79 20 53 the SecureWay S
0020: 65 63 75 72 69 74 79 20 53 65 72 76 65 72 20 66 ecurity Server f
0030: 6F 72 20 7A 2F 4F 53 20 28 52 41 43 46 29 or z/OS (RACF)
-[2]: ObjectId: 2.5.29.14 Criticality=false
SubjectKeyIdentifier [
KeyIdentifier [
0000: 05 6A CD 7F AE AF 89 78 99 A8 F1 5B 64 8B 9F AF .j.....x...[d...
0010: 73 1B 58 65 s.Xe
]
]
0[3]: ObjectId: 2.5.29.35 Criticality=false
AuthorityKeyIdentifier [
KeyIdentifier [
0000: 7E D1 7B 17 74 D3 AD D1 7D D8 F8 33 85 19 04 F8t.....3....
0010: 36 51 57 16 6QW.
]
]
0]
0]
Algorithm: [SHA1withRSA]
Signature:
0000: 73 0D FC E1 8A B3 42 E1 04 73 72 B1 C6 C9 87 54 s.....B..sr....T
0010: 87 57 02 FA 41 32 D8 B0 39 09 86 CB 6B 03 B6 F9 .W..A2..9...k...
0020: 62 8D 95 36 56 0E D4 D2 F7 7A 8D 4B FB 0B FD 91 b..6V....z.K....
0030: 89 A8 08 41 30 E2 27 DC 15 5F 2C F4 CD 2F 6B 8E ...A0.'..._/k.
0040: 21 2A 88 53 46 27 68 9B 55 14 38 8E 1F 50 95 BC !*.SF'h.U.8..P..
0050: A8 46 F6 68 97 9E 7B 65 9E E8 A7 34 B2 C8 63 CF .F.h...e...4..c.
0060: 73 C8 4E 25 0A EF C5 8F 04 A4 EB 8C CC 33 84 26 s.N%.....3.&
0070: 5D FD 7C AD 7B 02 13 5A 86 A1 89 93 1E A4 93 63].....Z.....c
0]
X509TrustManagerImpl: Certificate [
[
 Version: V3
 Subject: CN=WAS CertAuth, C=US
 Signature Algorithm: SHA1withRSA, OID = 1.2.840.113549.1.1.5
0 Key: IBMJCE RSA Public Key:
modulus:
1167408593733331602218385578183389496484587418638676352829560040529918
40558681208199977833401609895748222369066230329785148883251144

```

```

2382911186804921983976695395381692334250582278359056431484427844566504
41491799952592864895242987037929408453455627552772317382077015
828713585220212502839546496071839496308430393
public exponent:
65537
0 Validity: [From: Fri Jul 25 05:00:00 GMT 2003,
 To: Sat Jul 24 04:59:59 GMT 2010]
 Issuer: CN=WAS CertAuth, C=US
 SerialNumber: [0]
0Certificate Extensions: 4
 [1]: ObjectId: 2.16.840.1.113730.1.13 Criticality=false
 Extension unknown: DER encoded OCTET string =
 0000: 04 3C 13 3A 47 65 6E 65 72 61 74 65 64 20 62 79 .<.:Generated by
 0010: 20 74 68 65 20 53 65 63 75 72 65 57 61 79 20 53 the SecureWay S
 0020: 65 63 75 72 69 74 79 20 53 65 72 76 65 72 20 66 ecurity Server f
 0030: 6F 72 20 7A 2F 4F 53 20 28 52 41 43 46 29 or z/OS (RACF)
-[2]: ObjectId: 2.5.29.14 Criticality=false
 SubjectKeyIdentifier [
 KeyIdentifier [
 0000: 7E D1 7B 17 74 D3 AD D1 7D D8 F8 33 85 19 04 F8t.....3....
 0010: 36 51 57 16 6QW.
]
]
0[3]: ObjectId: 2.5.29.15 Criticality=true
 KeyUsage [
 Key_CertSign
 CrI_Sign
]
0[4]: ObjectId: 2.5.29.19 Criticality=true
 BasicConstraints:[
 CA:true
 PathLen:2147483647
]
0]
 Algorithm: [SHA1withRSA]
 Signature:
 0000: 43 88 AB 19 5D 00 54 57 5E 96 FA 85 CE 88 4A BF C...].TW^.....J.
 0010: 6E CB 89 4C 56 BE EF E6 8D 2D 74 B5 83 1A EF 9C n..LV....-t....
 0020: B3 82 F2 16 84 FA 5C 50 53 2A B4 FD EB 27 98 5D\'PS*...'.]
 0030: 43 48 D3 74 85 21 D1 E1 F2 63 9E FB 58 2A F3 6A CH.t.!...c..X*.j
 0040: 44 D2 F5 7D B2 55 B9 5E 32 11 78 B6 34 8E 4B 1D D....U.^2.x.4.K.
 0050: F3 82 1D C1 5F 7B 3F AD C9 29 FA FF D1 D1 13 2C?..).....
 0060: 57 F7 7B 51 02 99 6F ED 54 E1 51 34 B8 51 BE 97 W..Q..o.T.Q4.Q..
 0070: 30 AC 4F 89 AB AA 8A B2 E1 40 89 2E 18 C7 0E 15 0.O.....@.....
0]
 JSSEContext: handleConnection[Socket[addr=boss0106.plex1.l2.ibm.com
/9.38.48.108,port=9443,localport=2167]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2168]]

 JSSEContext: handleConnection[Socket[addr=boss0106.plex1.l2.ibm.com
/9.38.48.108,port=2235,localport=8878]]
 JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8879,localport=2236]]
 JSSEContext: handleSession[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8880,localport=2238]]
 JSSEContext: confirmPeerCertificate[Socket
[addr=BOSSXXXX.PLEX1.L2.IBM.COM
/9.38.48.108,port=8880,localport=2238]]
 X509TrustManagerImpl: checkServerTrusted
 X509TrustManagerImpl: Certificate [

 [
 Version: V3
 Subject: CN=jservice, OU=SWG, O=IBM, C=US
 Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4
0 Key: IBMJCE RSA Public Key:

```

```

modulus:
100949966922395090747968287561185391075683695663138899555389506686622953
008589748001058216362638201577071902071311277365773252660799
1287811829472738023126999835565278786157922922449953171124365624914
89904381884265119355037731265408654007388863303101746314438337601
264540735679944205391693242921331551342247891
public exponent:
65537
0 Validity: [From: Fri Jun 21 20:08:18 GMT 2002,
 To: Thu Mar 17 20:08:18 GMT 2005]
 Issuer: CN=jserver, OU=SWG, O=IBM, C=US
 SerialNumber: [3d1387b2]
0]
 Algorithm: [MD5withRSA]
 Signature:
0000: 54 DC B5 FA 64 C9 CD FE B3 EF 15 22 3D D0 20 31 T...d....."=. 1
0010: 99 F7 A7 86 F9 4C 82 9F 6E 4B 7B 47 18 2E C6 25L..nK.G...%
0020: 5B B2 9B 78 D8 76 5C 82 07 95 DD B8 44 62 02 62 [...x.v\.....Db.b
0030: 60 2A 0A 6D 4F B9 0A 98 14 27 E9 BB 1A 84 8A D1 ~*.m0.....'.....
0040: C2 22 AF 70 9E A5 DF A2 FD 57 37 CE 3A 63 1B EB ".p.....W7.:c..
0050: E8 91 98 9D 7B 21 4A B5 2C 94 FC A9 30 C2 74 72!J.,...0.tr
0060: 95 01 54 B1 29 E7 F8 9E 6D F3 B5 D7 B7 D2 9E 9B ..T.)...m.....
0070: 85 D8 E4 CF C2 D5 3B 64 F0 07 17 9E 1E B9 2F 79;d...../y
0]
X509TrustManagerImpl: Certificate [
[
 Version: V3
 Subject: CN=jserver, OU=SWG, O=IBM, C=US
 Signature Algorithm: MD5withRSA, OID = 1.2.840.113549.1.1.4
0 Key: IBMJCE RSA Public Key:
modulus:
100949966922395090747968287561185391075683695663138899555389506
686622953008589748001058216362638201577071902071311277365773252660799
12878118294727380231269998355652787861579229224499531711243656249
1489904381884265119355037731265408654007388863303101746314438337601
264540735679944205391693242921331551342247891
public exponent:
65537
0 Validity: [From: Fri Jun 21 20:08:18 GMT 2002,
 To: Thu Mar 17 20:08:18 GMT 2005]
 Issuer: CN=jserver, OU=SWG, O=IBM, C=US
 SerialNumber: [3d1387b2]
0]
 Algorithm: [MD5withRSA]
 Signature:
0000: 54 DC B5 FA 64 C9 CD FE B3 EF 15 22 3D D0 20 31 T...d....."=. 1
0010: 99 F7 A7 86 F9 4C 82 9F 6E 4B 7B 47 18 2E C6 25L..nK.G...%
0020: 5B B2 9B 78 D8 76 5C 82 07 95 DD B8 44 62 02 62 [...x.v\.....Db.b
0030: 60 2A 0A 6D 4F B9 0A 98 14 27 E9 BB 1A 84 8A D1 ~*.m0.....'.....
0040: C2 22 AF 70 9E A5 DF A2 FD 57 37 CE 3A 63 1B EB ".p.....W7.:c..
0050: E8 91 98 9D 7B 21 4A B5 2C 94 FC A9 30 C2 74 72!J.,...0.tr
0060: 95 01 54 B1 29 E7 F8 9E 6D F3 B5 D7 B7 D2 9E 9B ..T.)...m.....
0070: 85 D8 E4 CF C2 D5 3B 64 F0 07 17 9E 1E B9 2F 79;d...../y
0]
JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM/
9.38.48.108,port=8880,localport=2238]]
JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM/
9.38.48.108,port=8880,localport=2239]]
JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM/
9.38.48.108,port=8880,localport=2240]]
JSSEContext: handleConnection[Socket[addr=BOSSXXXX.PLEX1.L2.IBM.COM/
9.38.48.108,port=8880,localport=2241]]

```

## Tracing security

The classes which implement WebSphere Application Server security are:

- com.ibm.ws.security.\*
- com.ibm.websphere.security.\*
- com.ibm.WebSphereSecurityImpl.\*

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

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

---

## JSP processor troubleshooting tips

If you are having difficulty using the JSP engine, try these steps:

1. Determine whether other resources such as .html files or servlets are being requested and displayed correctly. If they are not, the problem probably lies at a deeper level, such as with the HTTP server.
2. If other resources are being displayed correctly, determine whether the JSP processor has started normally:
  - Browse the logs of the server hosting the JSP files you are trying to access. The following messages indicate that the JSP processor has started normally:  
Extension Processor [class com.ibm.ws.jsp.webcontainerext.JSPExtensionProcessor] was initialized successfully.  
Extension Processor [class com.ibm.ws.jsp.webcontainerext.JSPExtensionProcessor] has been associated with patterns [\*.jsp \*.jspx \*.jsw \*.jsw ].

If the JSP processor fails to load, you will see a message such as

```
No Extension Processor found for handling JSPs.
JSP Processor not defined. Skipping : jspfilename.
```

in the server log files.

3. If the JSP engine has started normally, the problem may be with the JSP file itself.
  - The JSP may have invalid JSP syntax and could not be processed by the JSP Processor. Examine the server log files of the target application for invalid JSP directive syntax messages. Errors similar to the following in a browser indicate this kind of problem:

```
Message: /filename.jsp(2,1)JSPG0076E: Missing required attribute page
for jsp element jsp:include
```

This example indicates that line 2, column 1 of the named JavaServer Pages file is missing a mandatory attribute for the jsp:include action. Similar messages are displayed for other syntax errors.

- Examine the target application server's server log files for problems with invalid Java syntax. Errors similar to **Message: Unable to compile class for JSP** in a browser indicate this kind of problem.

The error message output from the Javac compiler will be found in the server log files. It might look like:

```
JSPG0091E: An error occurred at line: 2 in the file: /myJsp.jsp
JSPG0093E: Generated servlet error: c:\WASROOT\temp\ ...
test.war\myJsp.java:16: myInt is already defined in com.ibm.ws.jsp20._myJsp
int myInt = 122; String myString = "number is 122"; static int myStaticInt=22;
int myInt=121;
 ^ 1 error
```

Correct the error in the JSP file and retry the file.

If none of these steps solves the problem, check to see if the problem is identified and documented using the links in Diagnosing and fixing problems: Resources for learning.

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

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

---

## Object request broker component troubleshooting tips

### Enabling tracing for the Object Request Broker component

The object request broker (ORB) service is one of the WebSphere Application Server run time services. Tracing of messages sent and received by the ORB is a useful starting point for troubleshooting the ORB service. You can selectively enable or disable tracing of ORB messages for each server in a WebSphere Application Server installation, and for each application client.

For instructions on how to set trace controls so that tracing occurs for the ORB subcomponent, see Setting trace controls for IBM service.

### Log files and messages associated with Object Request Broker

For a summary of how messages get routed in WebSphere Application Server, see Message routing.

### Java packages containing the Object Request Broker service

The ORB service resides in the following Java packages:

- com.ibm.com.CORBA.\*
- com.ibm.rmi.\*
- com.ibm.ws.orb.\*
- com.ibm.ws.orbimpl.\*
- com.ibm.ws390.orb.\*
- org.omg.CORBA.\*
- javax.rmi.CORBA.\*

JAR files that contain the previously mentioned packages include:

- *java\_install\_dir/jre/lib/ext/ibmorb.jar*
- *java\_install\_dir/jre/lib/ext/iwsorbutil.jar*
- *install\_dir/lib/iwsorb.jar*
- *install\_dir/lib/runtimews390.jar*
- *install\_dir/lib/bootstrap.jar*

### Tools used with Object Request Broker

The tools used to compile Java remote interfaces to generate language bindings used by the ORB at runtime reside in the following Java packages:

- com.ibm.tools.rmic.\*
- com.ibm.idl.\*

The JAR file that contains the packages is *install\_dir/java/lib/ibmtools.jar*.

## Object Request Broker properties

The ORB service requires a number of ORB properties for correct operation. It is not necessary for most users to modify these properties, and it is recommended that only your system administrator modify them when required. Consult IBM Support personnel for assistance. The properties reside in the `orb.properties` file, located at `install_dir/java/jre/lib/orb.properties`.

If none of these steps fixes your problem, check to see if the problem has been identified and documented by looking at the available online support (hints and tips, technotes, and fixes).

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

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

---

## Web services security troubleshooting tips

Troubleshooting Web services security is best done by reviewing the configurations in the Assembly Toolkit so that you can match up the client and server request and the response configurations. These configurations must match. A client request sender configuration must match a server request receiver configuration. For encryption to successfully occur, the public key of the receiver must be exported to the sender and this key must be configured properly in the encryption information. For authentication, you must specify the method used by the client in the login mapping of the server. Also, you must correctly specify the actor URI at each point in the configuration with the same URI string. The following includes a list of generic troubleshooting steps that you can perform. A listing of specific symptoms and solutions is provided after these steps.

### Steps for this task:

1. Verify that the client security extensions and server security extensions match on each downstream call for the following senders and receivers:
  - Request sender and request receiver
  - Response sender and response receiver
2. Verify that when the Add Created Time Stamp option is enabled on the client-side that the server has the Add Received Time Stamp option configured. You must configure the security extensions in the Assembly Toolkit
3. Verify that the client security bindings and the server security bindings are correctly configured. When the client authentication method is signature, make sure that the server has a login mapping. When the client uses the public key `cn=Bob,o=IBM,c=US` to encrypt the body, verify that this Subject is a personal certificate in the server key store so that it can decrypt the body with the private key. You can configure the security bindings using either the Assembly Toolkit or the WebSphere Application Server administrative console.
4. Check the `SystemOut.log` file in the `${USER_INSTALL_ROOT}/logs/server1` directory (`server1` changes depending upon the server name) for messages that might provide information about the problem.
5. Enable trace for Web services security by using the following trace specification:  
`com.ibm.xml.soapsec.*=all=enabled:com.ibm.ws.webservices.*=all=enabled:  
com.ibm.wsspi.wssecurity.*=all=enabled:com.ibm.ws.security.*=all=enabled:  
SASRas=all=enabled`

Type the previous three lines as one continuous line.

- Specific symptoms:

**Symptom:** WSEC5061E: The SOAP Body is not signed

**Solution:** This error usually occurs whenever the SOAP security handler does not load properly, and does not sign the SOAP body not to be signed. The SOAP security handler is typically the first validation that occurs on the server-side, so a multitude of problems can cause this message to display. The error might be caused by invalid actor URI configurations. You can configure the actor Universal Resource Identifier (URI) at the following locations within the Assembly Toolkit:

From the Web services client editor within the Assembly Toolkit for client configurations:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **ActorURI** field.
- Click **Security Extensions > Request Sender Configuration section > Details** and indicate the actor information in the **Actor** field.

From the Web Services Editor within the Assembly Toolkit for server configurations:

- Click **Security Extensions > Server Service Configuration** section. Verify that the actor URI has the same actor string as the client-side.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

The actor information on both the client and the server must refer to the same string. When the actor fields on the client and the server match, the request or response is acted upon instead of being forwarded downstream. The actor fields might be different when you have Web services acting as a gateway to other Web services. However, in all other cases, verify that the actor information matches on the client and server. When the Web services implementation is acting as a gateway and it does not have the same actor configured as the request passing through the gateway, this Web services implementation does not process the message from the client. Instead, it sends the request downstream. The downstream process that contains the correct actor string processes the request. The same situation occurs for the response. Therefore, it is important that you verify that the appropriate client and server actor fields are synchronized.

Additionally, the error can appear when you do not specify that the body is signed in the client configuration. To sign the body part of the message using the Web service client editor in the Assembly Toolkit, click **Security Extensions > Request Sender Configuration > Integrity** and select the message parts to sign.

- **Symptom:** WSEC5075E: No security token found that satisfies any one of the authentication methods.

**Solution:** Verify that the client and server login configuration information matches in the security extensions. Also, verify that the client has a valid login binding and that the server has a valid login mapping in the security bindings. You can check this information by looking at the following locations in the Assembly Toolkit:

From the Web services client editor within the Assembly Toolkit for client configurations:

- Click **Security Extensions > Request Sender Configuration > Login Configuration** verify the authentication method.



- Click **Port Binding > Security Request Sender Binding Configuration > Login Binding** verify the authentication method and other parameters.

From the Web Services Editor within the Assembly Toolkit for server configurations:

- Click **Security Extensions > Request Receiver Service Configuration Details > Login Configuration** and verify the authentication method.
- Click **Binding Configurations > Request Receiver Binding Configuration Details > Login Mapping** and verify the authentication method and other parameters.

Also, make sure that the actor URI specified on the client and server matches. You can configure the actor URI at the following locations within the Assembly Toolkit:

From the Web services client editor within the Assembly Toolkit for client configurations:

- Click **Security Extensions > Client Service Configuration Details** and indicate the actor information in the **ActorURI** field.
- Click **Security Extensions > Request Sender Configuration section > Details** and indicate the actor information in the **Actor** field.

From the Web services editor within the Assembly Toolkit for server configurations:

- Click **Security Extensions > Server Service Configuration** section. Make sure that the **Actor URI** field has the same actor string as the client side.
- Click **Security Extensions > Response Sender Service Configuration Details > Details** and indicate the actor information in the **Actor** field.

- **Symptom:** WSEC5094E: No UsernameToken of trusted user was found or the login failed for the user while the TrustMode is BasicAuth.

**Solution:** This situation occurs when you have IDAssertion configured in the login configuration as the authentication method. On the sending Web service, configure a trusted basic authentication entry in the login binding. Then, on the server side, verify that the trusted ID evaluator has a property set that contains the user name of this basic authentication entry. To configure the client for identity assertion, see:

- Configuring the client for identity assertion authentication: Specifying the method
- Configuring the client for identity assertion authentication: Collecting the authentication information

To configure the server for identity assertion, see:

- Configuring the server to handle identity assertion authentication
- Configuring the server to validate identity assertion authentication information

- **Symptom:** The following authorization error occurs with UNAUTHENTICATED as the security name: SECJ0053E: Authorization failed for /UNAUTHENTICATED while invoking (Home)com/ibm/wssvt/tc/pli/ejb/Beneficiary findBeneficiaryBySsNo(java.lang.String):2 securityName: /UNAUTHENTICATED;accessID: null is not granted any of the required roles: AgentRole

**Solution:** This situation occurs because a login configuration is not being configured or Web services Security is not configured from a client to a server. When the request arrives at the server and authentication information is not received, the UNAUTHENTICATED user is set on the thread. Authorization

returns this error if there are any roles assigned to the resource except for the special "Everyone" role, which supports access by anyone.

If the client successfully authenticates to an EJB file but the EJB file calls a downstream Enterprise JavaBean (EJB) file that is not configured with Web services security or transport security, such as HTTP user ID and password, an error can occur for this downstream request. Using the Assembly Toolkit, verify that the enterprise archive (EAR) file for both client and server has the correct security extensions and security bindings. For more information, see:

- Configuring the client security bindings using the Assembly Toolkit
- Configuring the security bindings on a server acting as a client using the administrative console
- Configuring the server security bindings using the Assembly Toolkit
- Configuring the server security bindings using the administrative console

---

## Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, except those expressly designated by IBM, is the user's responsibility.

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

IBM Director of Licensing  
IBM Corporation  
500 Columbus Avenue  
Thornwood, New York 10594 USA



---

## Trademarks and service marks

The following terms are trademarks of IBM Corporation in the United States, other countries, or both:

- AIX
- AS/400
- CICS
- Cloudscape
- DB2
- DFSMS
- Domino
- Everyplace
- iSeries
- IBM
- IMS
- Informix
- iSeries
- Language Environment
- Lotus
- MQSeries
- MVS
- OS/390
- RACF
- Redbooks
- RMF
- SecureWay
- SupportPac
- Tivoli
- ViaVoice
- VisualAge
- VTAM
- WebSphere
- z/OS
- zSeries

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

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

UNIX is a registered trademark of The Open Group in the United States and other countries.

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