

IBM WebSphere Application Server - Express for IBM i,
Version 8.5

Troubleshooting and support



Note

Before using this information, be sure to read the general information under “Notices” on page 229.

Compilation date: June 7, 2012

© Copyright IBM Corporation 2012.

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

Contents

How to send your comments	vii
Using this PDF	ix
Chapter 1. How do I troubleshoot?	1
Chapter 2. Liberty profile: Troubleshooting tips	3
Liberty profile: Trace and logging	6
Liberty profile: Runtime environment known restrictions	7
Liberty profile: Messages	9
Chapter 3. Debugging applications	13
Debugging components in the IBM Rational Application Developer for WebSphere	14
Debugging Service details	15
Enable service at server startup	15
JVM debug port	15
JVM debug arguments	15
Debug class filters	15
Chapter 4. Adding logging and tracing to your application	17
Using Java logging in an application	18
Using a logger	19
Java logging	28
Configuring the logger hierarchy	29
Creating log resource bundles and message files.	30
Logger.properties file for configuring logger settings	32
Configuring applications to use Jakarta Commons Logging	33
Jakarta Commons Logging	34
Configurations for the WebSphere Application Server logger.	37
Programming with the JRas framework	39
JRas logging toolkit.	39
JRas Extensions	41
JRas messages and trace event types.	49
Instrumenting an application with JRas extensions	52
Logging Common Base Events in WebSphere Application Server.	58
The Common Base Event in WebSphere Application Server.	59
Logging with Common Base Event API and the Java logging API	72
java.util.logging -- Java logging programming interface	82
Logger.properties file	83
Logging Common Base Events in WebSphere Application Server.	84
Showlog commands for Common Base Events	84
Chapter 5. Configuring Java logging using the administrative console	85
Log streams and expected output	86
Log level settings	89
Changing the message IDs used in log files.	92
Converting log files to use IBM unique Message IDs	93
convertlog command	93
MessageConverter class	94
HTTP error, FRCA, and NCSA access log settings	94
Enable logging service at server start-up	95
NCSA access logging	95
Error logging	96

Chapter 6. Using High Performance Extensible Logging to troubleshoot applications	97
High Performance Extensible Logging (HPEL)	97
Basic mode and HPEL mode	102
Changing from basic mode to HPEL logging and tracing	104
Changing from HPEL to basic mode logging and tracing	105
Determining which of basic mode and HPEL mode is enabled	107
Configuring HPEL with wsadmin scripting	108
Configuring HPEL	111
HPEL logging and trace settings	111
HPEL log configuration settings	111
HPEL trace configuration settings	113
HPEL text log configuration settings	115
Change log and trace mode settings	117
Log viewer settings	117
Log view table	117
Content and filtering details	118
Server instance	118
View contents	118
System Out	118
System Error	119
Logs and trace	119
Filtering	119
Include loggers	119
Exclude loggers	119
Message contents	119
Event timing	119
From	119
On (first occurrence)	120
Until	120
On (second occurrence)	120
LogViewer command-line tool	120
Monitoring application logging using JMX notifications	123
Chapter 7. Using Cross-Component Trace to troubleshoot applications	127
Cross Component Trace (XCT)	128
Configuring XCT with wsadmin scripting	130
Chapter 8. Using sensitive log and trace guard	133
Sensitive log and trace guard	133
Enabling and disabling sensitive log and trace guard	134
Maintaining sensitive log and trace guard lists	134
Chapter 9. Diagnosing problems (using diagnosis tools)	137
Chapter 10. Using basic or traditional message logs to troubleshoot applications	139
Viewing JVM logs	140
JVM log interpretation	141
Configuring the JVM logs	143
Java virtual machine (JVM) log settings	143
Monitoring application logging using JMX notifications	145
Process logs	147
Configuring the service log	148
IBM service log settings	148
Viewing the service log	149
Chapter 11. Working with trace	151

Enabling trace on client and stand-alone applications	151
Enabling trace at server startup	152
Enabling trace on a running server	153
Managing the application server trace service	154
Diagnostic trace service settings	155
Trace Output.	155
Trace Output Format.	156
Runtime tab	156
Trace Output.	156
Select a server to configure logging and tracing	157
Server	157
Host name	157
Version.	157
Type.	157
Status	158
Log and trace settings	158
Switch to HPEL Mode	158
Diagnostic Trace	158
Java virtual machine (JVM) Logs	158
Process Logs	158
IBM Service Logs	159
Change Log Level Details	159
NCSA access and HTTP error logging	159
Chapter 12. Troubleshooting class loaders	161
Class loading exceptions	163
osgiCfgInit script	168
Class loader viewer service settings	168
Enable service at server startup.	169
Enterprise application topology	169
Enterprise applications topology.	169
Class loader viewer settings	169
Class Loader	170
Search settings.	171
Search type	171
Search terms	171
Chapter 13. Configuring the hang detection policy	173
Hung threads in Java Platform, Enterprise Edition applications	175
Example: Adjusting the thread monitor to affect server hang detection.	177
Chapter 14. Working with troubleshooting tools	179
Gathering information with the collector tool (deprecated)	179
Collector tool output	181
collector command - summary option.	181
First failure data capture (FFDC)	182
Configuring first failure data capture log file purges	183
Chapter 15. Working with Diagnostic Providers	185
Diagnostic Providers	185
Diagnostic Provider IDs.	187
Diagnostic Provider configuration dumps, state dumps, and self tests	188
Diagnostic Provider registered attributes and registered tests	189
Diagnostic Provider names	191
The simpler interfaces provided by the Diagnostic Service MBean	191
Creating a Diagnostic Provider	191

Diagnostic Provider Extensible Markup Language	192
Choosing a Diagnostic Provider name	193
Implementing a Diagnostic Provider	194
Creating a Diagnostic Provider registration XML file	200
Associating a Diagnostic Provider ID with a logger	200
Static Assignment	200
Dynamic Assignment.	201
Using Diagnostic Providers from wsadmin scripts	201
Viewing the run time configuration of a component using Diagnostic Providers	203
Configuration data quick link or server selection	203
Diagnostic Providers (selection).	203
Configuration data.	204
Viewing the run time state data or configuring the state data collection specifications for a Diagnostic Provider	204
Diagnostic Provider State Collection Specification	205
State Data Quick Link or Server Selection	206
State data.	206
Detailed state specification	207
Change state specification.	207
Modifying the State Collection Specification from wsadmin scripts	208
Running a self diagnostic on a Diagnostic Provider	208
Tests Quick Link or Server Selection	209
Test selection	209
Test Results	209
Test result details	210
Chapter 16. Troubleshooting help from IBM	211
Diagnosing and fixing problems: Resources for learning	211
Using IBM Support Assistant	212
Using the IBM Support Assistant Data Collector.	213
Chapter 17. Default behavior for OutOfMemory exceptions	215
Chapter 18. Configuring the memory leak policy	217
Memory leaks in Java Platform, Enterprise Edition applications	222
Chapter 19. Collecting Java dumps and core files using the administrative console	225
Java dump and core collection	225
Chapter 20. Directory conventions	227
Notices	229
Trademarks and service marks	231
Index	233

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 email form appears.
 3. Fill out the email form as instructed, and submit your feedback.
- To send comments on PDF books, you can email your comments to: **wasdoc@us.ibm.com**.

Your comment should pertain to specific errors or omissions, accuracy, organization, subject matter, or completeness of this book. 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.

For technical questions and information about products and prices, please contact your IBM branch office, your IBM business partner, or your authorized remarketer. When you send comments to IBM, you grant IBM a nonexclusive right to use or distribute your comments in any way it believes appropriate without incurring any obligation to you. IBM or any other organizations will only use the personal information that you supply to contact you about your comments.

Using this PDF

Links

Because the content within this PDF is designed for an online information center deliverable, you might experience broken links. You can expect the following link behavior within this PDF:

- Links to Web addresses beginning with `http://` work.
- Links that refer to specific page numbers within the same PDF book work.
- The remaining links will *not* work. You receive an error message when you click them.

Print sections directly from the information center navigation

PDF books are provided as a convenience format for easy printing, reading, and offline use. The information center is the official delivery format for IBM WebSphere Application Server documentation. If you use the PDF books primarily for convenient printing, it is now easier to print various parts of the information center as needed, quickly and directly from the information center navigation tree.

To print a section of the information center navigation:

1. Hover your cursor over an entry in the information center navigation until the **Open Quick Menu** icon is displayed beside the entry.
2. Right-click the icon to display a menu for printing or searching your selected section of the navigation tree.
3. If you select **Print this topic and subtopics** from the menu, the selected section is launched in a separate browser window as one HTML file. The HTML file includes each of the topics in the section, with a table of contents at the top.
4. Print the HTML file.

For performance reasons, the number of topics you can print at one time is limited. You are notified if your selection contains too many topics. If the current limit is too restrictive, use the feedback link to suggest a preferable limit. The feedback link is available at the end of most information center pages.

Chapter 1. How do I troubleshoot?

Follow these shortcuts to get started quickly with popular tasks.

When you visit a task in the information center, look for the **IBM® Suggests** feature at the bottom of the page. Use it to find available tutorials, demonstrations, presentations, developerWorks® articles, Redbooks®, support documents, and more.

Add tracing and logging to your applications

* For more detailed information on enabling traces by using scripting, see the Troubleshooting with scripting chapter in the *Administering applications and their environment* PDF book.

Debug WebSphere® applications during development

Detect hung threads

Detect product configuration file problems

Set traces and logs with the console

Set traces and logs with scripting*

Work with message logs

Using IBM Support Assistant

Using HPEL to troubleshoot applications

Using JSR47 for logging

Using JSR47 for logging: Configuring access logs

Chapter 2. Liberty profile: Troubleshooting tips

Tips for troubleshooting the Liberty profile.

To help you identify and resolve problems, the product has a unified logging component. See “Liberty profile: Trace and logging” on page 6.

If you receive an exception message, information about the message is available in “Liberty profile: Messages” on page 9.

The Java API document for each Liberty profile API is detailed in the Programming Interfaces (APIs) section of the information center, and is also available as a JAR file under the `/dev/ibm-api/javadoc` directory of the server image.

IBM i For details of the main known restrictions that apply when using the Liberty profile, see “Liberty profile: Runtime environment known restrictions” on page 7.

Here is a set of tips to help you troubleshoot commonly experienced problems:

- Troubleshooting JDKs
- “Troubleshooting security”
- “Troubleshooting LDAP” on page 5
- “Troubleshooting SSL” on page 5

Check that your JDKs are at a supported level

If you experience problems that are not readily explained, check that the JDKs you are using are compliant with Java Version 6 or later, and are at a current service level. See “Minimum supported Java levels” on page 7.

Troubleshooting security

This section describes some common security problems and solutions you can choose.

SESN008E: A user authenticated as anonymous has attempted to access a session owned by user:LdapRegistry/cn=steven,o=myCompany,c=US.

This exception happens when an unauthenticated user tries to access a session created by an authenticated user. Session security which is enabled by default prevents unauthorized access of the sessions. Only the user who created a session can access it. See session security for more information.

This exception can happen when you use a JSP (`login.jsp` for example) for your form-login and the SSO token sent by the browser is expired. Because the SSO token is expired, the user is prompted to log in again using the `login.jsp` page configured for the form-login. The jsp page, by default, tries to get a session. This session was originally created by the user whose token is expired. However, since the token is expired the authentication of the user did not succeed and hence there are no credentials established when accessing this session resulting in this exception. To avoid this exception, either restart the browser which starts a new session, or configure the `login.jsp` file to not create the session by default. If you choose to update the `login.jsp` file, set `<%@ page session="false" %>`.

CWWKS9104A: Authorization failed for user {0} while invoking {1} on {2}. The user is not granted access to any of the required roles: {3}.

This error occurs when the user does not have authorization to the roles required by the application. Make sure that the user or the group it belongs to is mapped to at least one of the roles mentioned in the error message. Recall that a user-to-role mapping defined in the

ibm-application-bnd.xmi/xml file takes precedence over a mapping defined in the server.xml file. Check both resources to ensure that the correct mapping is defined. See Configuring authorization for applications on the Liberty profile.

CWWKZ0013E: It is not possible to start two applications called {0} followed by unexpected security behavior and error messages such as CWWKS9104A.

This error occurs when you have specified your application in both the server.xml (using the application element) and in the dropins folder. As a result, the application is attempted to be installed twice and the security configuration present in the server.xml file might or might not take effect. To fix this, you must remove your application from the dropins folder and copy it to another directory. If you have to leave it in the dropins folder, then you must also disable application monitoring by setting the following in your server.xml file:

```
<applicationMonitor dropinsEnabled="false"/>
```

An unauthenticated user was able to access my servlet or JSP.

A user with a principal of **UNAUTHENTICATED** (or the unauthenticated SAF user on z/OS®) is created when authentication fails or when your servlet or JSP is unprotected. An unauthenticated user is able to access your servlet or JSP if you do not define any security constraints or if you map the special subject of **EVERYONE** to the role required by your application. Review the user-to-role mappings in both the ibm-application-bnd.xmi/xml and the server.xml files. Consider the following options:

- If your servlet or JSP is unprotected, then consider adding security constraints to the deployment descriptor of your application. See Liberty profile: Authentication.
- If you do not want any unauthenticated user to access your application, then consider removing the **EVERYONE** special subject from the mapping for that role. See Configuring authorization for applications on the Liberty profile.
- If you believe that a certain user cannot be authorized to your servlet or JSP, then review who is mapped to that role by examining both the ibm-application-bnd.xmi/xml and the server.xml files. You can map a role to a user, group, or special subject. If your role is mapped to the **EVERYONE** special subject, then any user is allowed access. If your role is mapped to the **ALL_AUTHENTICATED** special subject, then any authenticated user is allowed access. Remove these special subjects if you want to further limit who can access your servlet or JSP. Finally, check what group the user belongs to. Although the user might not have explicit access, the group might be mapped to the role. In this case, consider removing the user from the authorized group or considering removing the group from role mapping. See Configuring authorization for applications on the Liberty profile.

Single Sign On (SSO) does not work.

If SSO does not work, check that the different Liberty profile servers that use the same LTPA keys, password, and **ssoCookieName** attribute of webAppSecurity element, each server has the same Universal Time (UTC), and the same user registry is shared. Also, if the token expires or if the cookie is sent from a web browser after changing the user registry in a manner that is inconsistent, like modifying the realm or removing the user the cookie represents, the SSO fails and the user is prompted to enter the credential information again. See Customizing SSO configuration using LTPA cookies for the Liberty profile.

Debugging security public APIs.

WSSecurityHelper and RegistryHelper are loaded even if security is not enabled. For example, if a security feature - security-1.0 or zosSecurity-1.0 - is not specified. If security is not enabled, then WSSecurityHelper.isServerSecurityEnabled() and WSSecurityHelper.isGlobalSecurityEnabled() methods both return false, and RegistryHelper.getUserRegistry() method returns null.

Other security public API classes might not be loaded if security is not enabled. If you try to access these classes you might get a java.lang.NoClassDefFoundError exception if you try to call a method on the class.

To avoid getting `java.lang.NoClassDefFoundError` exceptions, you must first test to see whether security is enabled by calling the `WSSecurityHelper.isServerSecurityEnabled()` or `WSSecurityHelper.isGlobalSecurityEnabled()` class, and then call other security public API classes only when security is enabled. See Liberty profile: Security public APIs for examples of this coding technique.

Note: This behavior is different from the full profile. In full profile, all classes are always loaded regardless of whether security is enabled or not.

Troubleshooting LDAP

This section describes some common LDAP problems and solutions you can choose.

FFDC1015I: An FFDC Incident has been created: "javax.naming.ServiceUnavailableException: myldapserver.mycompany.com:636; socket closed com.ibm.ws.security.registry.ldap.internal.LdapRegistry 298

This message in `messages.log` indicates that the configured LDAP server cannot be reached. Check your LDAP server to verify that it is running and that its IP address can be accessed from the Liberty profile server.

The javax.naming.CommunicationException: simple bind failed: myldapserver.mycompany.com:636 [Root exception is javax.net.ssl.SSLHandshakeException: com.ibm.jsse2.util.g: PKIX path building failed: java.security.cert.CertPathBuilderException: unable to find valid certification path to requested target]

If you enable SSL on your LDAP server without copying the signer of the LDAP server into the truststore referenced in the `LDAPSSLSettings` element, a connection with the LDAP server fails with an SSL handshake error. Make sure that you copy the signer of the LDAP server into your truststore.

The javax.naming.CommunicationException: myldapserver.mycompany.com:389 [Root exception is java.net.BindException: Address already in use: connect]

This message might appear in the FFDC logs and indicates that the usable sockets available on the local server are exhausted, resulting in a failure when connecting to your LDAP server. Make sure the socket is not used and try again.

CWWKS1100A: Authentication did not succeed for user ID xxxxx. An invalid user ID or password was specified

This FFDC exception could happen on the server even though the user mentioned in the message above is a valid user on LDAP server under heavy load. LDAP configuration allows to add the property `reuseConnection=false` or disable it using the developer tools. To fix the problem, Setting this property to default value of `true`.

Troubleshooting SSL

This section describes some common SSL problems and solutions you can choose.

CWWKS9105E: Could not determine the SSL port for automatic redirection.

This error occurs when a user tries to access an application that redirects to an SSL port and the SSL port is not available. The port might not be available due to a missing SSL configuration or some error in the SSL configuration definition. Check the SSL configuration in the `server.xml` file to make sure that it exists and is correct. See Securing communications with the Liberty profile.

FFDC1015I: An FFDC Incident has been created: "java.lang.IllegalArgumentException: Unknown, incomplete configuration: missing id field com.ibm.ws.config.internal.cm.ManagedServiceFactoryTracker aSyncReadNupdate. Exception thrown while trying to read configuration and update ManagedServiceFactory. Exception = java.lang.IllegalArgumentException: Unknown, incomplete configuration: missing id field" at

ffdc_12.04.18_16.09.42.0.log

This error occurs when there is a keystore element in the configuration without an ID field. If you are using a minimal SSL configuration, set the ID field to `defaultKeyStore`.

Liberty profile: Trace and logging

The product has a unified logging component. It provides base implementations of trace and FFDC services, for runtime code and for your code, to gather debug information.

The trace and FFDC implementations both apply an initial configuration during static initialization. You can modify this default initial configuration by specifying properties in the server configuration files (see [Configuring the Liberty profile runtime environment](#)) or `bootstrap.properties` file (see [Specifying Liberty profile bootstrap properties](#)).

Messages are written to stdout as well as to the defined trace destination. OSGi logging output is intercepted and output through the trace support. There is also interception of `java.util.logging` output.

Trace configuration

The logging service can be controlled through the server configuration. Occasionally you need to set logging properties so they can take effect before the server configuration files are processed; in this case you set them in the `bootstrap.properties` file instead of the server configuration. You do not usually need to do this to get logging from your own code, which is loaded after server configuration processing, but you might need to do this to analyze problems in early server start or configuration processing.

Logging properties can be set in either `bootstrap.properties` or `server.xml` file. You use attributes in the `server.xml` file, or use equivalent properties in the `bootstrap.properties` file. Any settings in the `bootstrap.properties` file are used when the `server.xml` file is processed. If the logging properties in the `bootstrap.properties` file are not replaced or reset in the `server.xml` file, the property values in the `bootstrap.properties` file are used.

Table 1. Logging properties for the Liberty profile. Column 1 contains attributes that can be set in the `server.xml` file. Column 2 contains equivalent properties that can be used in the `bootstrap.properties` file. Column 3 provides description of each logging property.

Attribute	Equivalent property	Description
logDirectory	<code>com.ibm.ws.logging.log.directory</code>	This sets the directory for all log files, including FFDC.
maxFileSize	<code>com.ibm.ws.logging.max.file.size</code>	The maximum size a log file is allowed to reach before being rolled (in MB). The Liberty profile runtime only does size-based log rolling. To disable this, set the value to 0. The maximum file size is approximate.
maxFiles	<code>com.ibm.ws.logging.max.files</code>	If there is an enforced maximum file size, this setting is used to determine how many of each log file is kept. This setting is also applied to the number of exception logs (summaries of exceptions that occurred on any given day). So if this number is 10, you might have 10 message logs, 10 trace logs, and 10 exception summaries in the <code>ffdc/</code> directory.
consoleLogLevel	<code>com.ibm.ws.logging.console.log.level</code>	This filter controls the granularity of messages that go to the <code>console.log</code> file. The valid values are INFO, AUDIT, WARNING, ERROR and OFF. By default, the level is AUDIT.

Table 1. Logging properties for the Liberty profile (continued). Column 1 contains attributes that can be set in the `server.xml` file. Column 2 contains equivalent properties that can be used in the `bootstrap.properties` file. Column 3 provides description of each logging property.

Attribute	Equivalent property	Description
<code>messageFileName</code>	<code>com.ibm.ws.logging.message.file.name</code>	The message log has a default name of <code>message1.log</code> . This file always exists, and contains INFO and other (AUDIT, WARNING, ERROR, FAILURE) messages in addition to <code>System.out</code> and <code>System.err</code> . This log also contains timestamps and the issuing thread ID.
<code>traceFileName</code>	<code>com.ibm.ws.logging.trace.file.name</code>	The <code>trace.log</code> file is only created if additional or detailed trace is enabled. <code>stdout</code> is recognized as a special value, and causes trace to be directed to the original standard out stream.
<code>traceSpecification</code>	<code>com.ibm.ws.logging.trace.specification</code>	The trace string is used to selectively enabled trace. The default is <code>*=info=enabled</code> .
<code>traceFormat</code>	<code>com.ibm.ws.logging.trace.format</code>	This controls the format of the trace log. The default format for the Liberty profile is ENHANCED. You can also use BASIC and ADVANCED formats as in the full profile.

You can set logging properties in the server configuration file by selecting **Logging and Tracing** in the **Server Configuration** view in the developer tools, or by adding a logging element to the server configuration file like this:

```
<logging traceSpecification="*=audit=enabled:com.myco.mypackage.*=debug=enabled"/>
```

For the full logging configuration reference, including the detailed format of the `traceSpecification` property, see the logging element in Liberty profile: Configuration elements in the `server.xml` file.

Note: On all platforms, logs are written in the default system encoding.


Liberty profile: Runtime environment known restrictions

There are a small number of known restrictions that apply when working with the Liberty profile runtime environment.

Minimum supported Java levels

The minimum supported level for the JDK from Oracle is Java 6 update 26. For the Java SDK from IBM, the minimum supported level is 6.0 (J9 2.6) SR 1.

The JDK from IBM is available in many IBM products. For example, WebSphere Application Server Version 7 includes the Java 6 JDK from IBM.

Note:  To obtain the minimum supported Java level for IBM iSeries®, install IBM J2SE 6.0 32-bit JVM (5761-JV1 option 11) or IBM SE 6.0 64 bit (5761-JV1 option 12), and also install Java PTF group SF99562 level 19 (or later) for IBM i 6.1 or SF99572 level 8 (or later) for i 7.1

The install directory name and path cannot include non-ASCII characters

Recent JVMs do not fully support use of non-ASCII characters with the `-jar` and `-javaagent` commands. You should use only ASCII characters in your install directory names and paths.

Changing the JDBC data source at run time might result in JPA failures

If not specified through properties, OpenJPA detects and calculates the database dictionary type when the first entity manager is created and the database connection is made. This database dictionary type is used for all entity managers that are subsequently created. If the JDBC data source is changed while an application is running, the entity manager factory does not detect this change and continues to use the old dictionary for operations against the new data source. This can result in failures if the database is changed to a different vendor.

When you change a database to a different vendor, restart the application.

Modifying the dataSource, jdbcDriver, connectionManager and JDBC vendor properties at run time might result in JPA failures

If you update the configuration of dataSource, jdbcDriver, connectionManager or any of the JDBC vendor properties lists (for example, properties.db2.jcc or properties.oracle) while the server is running, you might see J2CA8040E failures. These failures state that multiple dataSource elements cannot be associated with a single connectionManager. These failures are generated even if your configuration associates only one connectionManager with the dataSource element.

When you make updates to the configuration of any of these JDBC resources, restart the server.

An application that relies on a result being returned by getRealPath must be deployed as an expanded application, not as a WAR file

The Java EE specification states that the getRealPath() method returns a null value if the content is being made available from a web archive (WAR) file. When you deploy a WAR file to the Liberty profile, the profile does not automatically extract the archive file into a directory structure. Therefore the application might fail to start. If your application relies on a result being returned by getRealPath(), you must deploy the application as an expanded web application, not as a WAR file. For example, you can manually extract the WAR file and copy the expanded application to the dropins directory.

WebSphere Application Server full profile scripts do not work with the Liberty profile

You cannot use any of the scripts under the bin directory of the WebSphere Application Server full profile to administer the Liberty profile.

Fileset restrictions

The following restrictions apply to Filesets:

- Filesets do not recursively explore subdirectories of the base directory. For example the following instructions are not supported:

```
<fileset id="testFileset" dir="\temp" includes="**\a.jar"/>
<fileset id="testFileset" dir="\temp" includes="a\a.jar"/>
<fileset id="testFileset" dir="\temp" includes="*\a.jar"/>
<fileset id="testFileset" dir="\temp" includes="a\b\a.jar"/>
```
- If you use symbolic links with Filesets, you must add a forward slash ("/") at the end of the dir attribute value. For example:

```
<fileset dir="{ihc.home}/" includes="*.jar"/>
```

beanvalidation-1.0 feature restrictions

For the beanvalidation-1.0 feature, the following restriction applies:

- There is no support for bean validation inside OSGi applications.

jaxrs-1.1 feature restriction

For the jaxrs-1.1 feature, Liberty profile does not support JAX-RS third-party APIs in the developer tools as in the full profile.

jpa-2.0 feature restrictions

For the jpa-2.0 feature, the following restrictions apply:

- There is no support for dynamic removal from the runtime.
- You cannot use Java 7 to compile and build your JPA applications. Applications that are built with Java 6 will run on the Liberty profile with Java 7.

jsf-2.0 feature restrictions

For the jsf-2.0 feature, the following restriction applies:

- There is no support for dynamic removal from the runtime.

jsp-2.2 feature restrictions

For the jsp-2.2 feature, the following restriction applies:

- There is no support for the useInMemory configuration option to only store the translated JSP file in memory.

Liberty profile: Messages

When you use the Liberty profile, you might encounter system messages. Each message has a unique message identifier, and includes an explanation of the problem, and details of any action that you can take to resolve the problem.

Liberty profile system messages are logged from a variety of sources, including application server components and applications. For the Liberty profile, the message identifier is 10 characters in length and has the following form:

CWXXX9999X

where:

CWXXX

is a five-character alphabetic message prefix that identifies the Liberty profile component.

9999 is a four-character numeric identifier used to identify the specific message for that component.

X is an optional alphabetic indicator that identifies the message type. (I=Informational, W=Warning, E=Error)

The Troubleshooter reference: Messages topic contains information about all WebSphere Application Server messages, indexed by message prefix. For each message there is an explanation of the problem, and details of any action that you can take to resolve the problem.

Table 2. Alphabetic message prefixes for Liberty profile messages

Liberty profile message prefix	Range	Liberty profile component
CWLIB	0001-0100	Transactions
CWLIB	0101-0200	z/OS Transaction Extensions
CWWKB		z/OS Native integration
CWWKB	0001-0050	z/OS Command Processing
CWWKB	0051-0100	z/OS Native code only

Table 2. Alphabetic message prefixes for Liberty profile messages (continued)

Liberty profile message prefix	Range	Liberty profile component
CWWKB	0101-0150	z/OS Core
CWWKB	0151-0200	z/OS WLM services
CWWKC	0000-0250	Annotation scanning
<i>CWWKE</i>		<i>Core kernel and kernel services</i>
CWWKE	0001-0099	Boot/Launcher
CWWKE	0100-0199	Service utilities
CWWKE	0200-0299	Location service
CWWKE	0300-0399	File install service
CWWKE	0400-0499	FileMonitor service
CWWKE	0500-0599	Extensible Class Scanner
CWWKF		Feature manager
CWWKG		Configuration manager
CWWKJ		Light touch administration
CWWKL	0001-0100	Class loading service
CWWKM	0001-0050	Artifact API messages (container factory)
CWWKM	0051-0100	Overlay Messages (all implementations)
CWWKM	0101-0150	Artifact API Zip implementation
CWWKM	0151-0200	Artifact API File implementation
CWWKM	0401-0450	Adaptable API Implementation messages. (adaptable module factory / adapter service)
CWWKM	1001-1100	Loose archive API (used by the Eclipse-based tools option that runs an application directly from the Eclipse workspace(The Eclipse platform supports a "virtual" application archive, where a set of files that appear to be in the same archive file are actually spread out across the Eclipse workspace. Liberty calls this a "loose archive", and the tools option that uses the loose archive API is called Run this application directly from the workspace)).
CWWKM	2000-2999	Ant and maven plugin
CWWKN	0001-0100	JNDI default namespace
<i>CWWKO</i>		<i>CFW components</i>
CWWKO	0000-0199	CFW
CWWKO	0200-0399	TCP
CWWKO	0400-0599	UDP
CWWKO	0600-0799	bytebuffer
CWWKO	0800-0899	SSL channel

Table 2. Alphabetic message prefixes for Liberty profile messages (continued)

Liberty profile message prefix	Range	Liberty profile component
CWWKO	0900-0999	SSL
<i>CWWKS</i>		<i>Security</i>
<i>CWWKS</i>	<i>0000 series</i>	<i>Security: General messages</i>
CWWKS	0000-0099	Security
CWWKS	0900-0999	Security quickStart
<i>CWWKS</i>	<i>1000 series</i>	<i>Security: Authentication services</i>
CWWKS	1000-1099	Authentication
CWWKS	1100-1199	JAAS authentication
CWWKS	1200-1299	TAI authentication
<i>CWWKS</i>	<i>2000 series</i>	<i>Security: Authorization services</i>
CWWKS	2000-2099	Authorization
CWWKS	2100-2199	Built-in authorization
CWWKS	2900-2999	SAF authorization
<i>CWWKS</i>	<i>3000 series</i>	<i>Security: Registry services</i>
CWWKS	3000-3099	Registry
CWWKS	3100-3199	Basic registry
CWWKS	3200-3299	LDAP registry
CWWKS	3300-3399	FileBased (VMM) registry
CWWKS	3800-3899	Custom registry
CWWKS	3900-3999	SAF registry
<i>CWWKS</i>	<i>4000 series</i>	<i>Security: Token services</i>
CWWKS	4000-4099	Token services
CWWKS	4100-4199	LTPA
CWWKS	4200-4299	Kerberos
CWWKS	4300-4399	SPNEGO
<i>CWWKS</i>	<i>9000 series</i>	<i>Security: Collaborators</i>
CWWKS	9100-9199	Web Collaborator (common code)
CWWKS	9200-9299	Web Application Collaborator
CWWKS	9300-9399	Web Administration Collaborator
CWWKT		HTTP transport/dispatcher
<i>CWWKX</i>		<i>JMX</i>
CWWKX	0001-0100	JMX Security
CWWKX	0101-0200	JMX REST Connector
CWWKX	0201-0300	JMX REST Client
<i>CWWKZ</i>		<i>Applications</i>
CWWKZ	0001-0100	Application manager
CWWKZ	0101-0200	WARs
CWWKZ	0201-0300	WABs
CWWKZ	0301-0400	EBA

Chapter 3. Debugging applications

To debug your application, you must use a development environment like the IBM Rational® Application Developer for WebSphere to create a Java project. You must then import the program that you want to debug into the project.

About this task

By following the steps below, you can import the WebSphere Application Server examples into a Java project. Two debugging styles are 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 do not need to 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.

Procedure

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 specify a name for the project, for example, WASExamples.
4. Click **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. Browse the directory for files.

 Go to the following directory: *profile_root/installedApps/node_name/DefaultApplication.ear/DefaultWebApplication.war*.

7. Select *DefaultWebApplication.war* in the left side of the Import dialog and then click **Finish**. This imports 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.

 Select **Properties** from the right-click menu. Choose the Java Build Path node and then select the Libraries tab. Click **Add External JARs** to add the following JAR files:

- *profile_root/installedApps/node_name/DefaultApplication.ear/Increment.jar*.

When you have added this JAR file, select it and use the **Attach Source** function to attach the *Increment.jar* file because it contains both the source and class files.

- *app_server_root/dev/JavaEE/j2ee.jar*
- *app_server_root/plugins/com.ibm.ws.runtime.jar*
- *app_server_root/plugins/com.ibm.ws.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. When 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 WASEExamples 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://your.server.name:9080/hitcount`

IBM i

What to do next

To learn more about debugging, launch The IBM Rational Application Developer for WebSphere, select **Help > Help Contents** and choose the **Debugger Guide bookshelf** entry. To learn about known limitations and problems that are associated with the IBM Rational Application Developer for WebSphere, see the IBM Rational Application Developer for WebSphere release notes. 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 Must gather documents page for information to gather to send to IBM Support.

Debugging components in the IBM Rational Application Developer for WebSphere

The IBM Rational Application Developer for WebSphere, included with the WebSphere Application Server on a separately-installable CD, includes debugging functionality that is built on the Eclipse workbench. Documentation for the IBM Rational Application Developer for WebSphere is provided with that product. To learn more about the debug components, launch the IBM Rational Application Developer for WebSphere, select **Help > Help Contents** and choose the **Developing > Debugging applications** bookshelf entries.

The IBM Rational Application Developer for WebSphere includes the following components:

The WebSphere Application Server debug adapter

which allows you to debug web objects that are running on WebSphere Application Server and that you have launched in a browser. These objects include enterprise beans, JavaServer Pages files, and servlets.

The JavaScript debug adapter

which enables server-side JavaScript debugging.

The Compiled language debugger

which allows you to detect and diagnose errors in compiled-language applications.

The Java development tools (JDT) debugger

which allows you to debug Java code.

All of the debug components in the IBM Rational Application Developer for WebSphere can be used for debugging locally and for remote debugging. To learn more about the debug components, launch the IBM Rational Application Developer for WebSphere, select **Help > Help Contents** and choose the **Developing > Debugging applications** bookshelf entries.

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 > Servers Types > WebSphere application servers > *server name* > Debugging service**.

You can enable a debug session on WebSphere Application Server on this page. Debugging can prove useful when your program behaves differently on the application server than on your local system.

Enable service at server 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.

Chapter 4. Adding logging and tracing to your application

You can add logging and tracing to applications to help analyze performance and diagnose problems in WebSphere Application Server.


About this task

Deprecation: The JRAs framework that is described in this information center is deprecated. However, you can achieve the same results using Java logging.

Designers and developers of applications that run with or under WebSphere Application Server, such as servlets, JavaServer Pages (JSP) files, enterprise beans, client applications, and their supporting classes, might find it useful to use Java logging for generating their application logging.

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

- Your messages are displayed in the WebSphere Application Server standard log files, using a standard message format with additional data, such as a date and time stamp that are added automatically.
- You can more easily correlate problems and events in your own application to problems and events that are associated with WebSphere Application Server components.
- You can take advantage of the WebSphere Application Server log file management features.



Procedure

1. Enable and configure any of the supported types of logging as needed. Use one of the following methods:
 - Configuring Java logging using the administrative console
 - Configuring applications to use Jakarta Commons Logging
2. Customize the properties to meet your logging needs. For example, enable or disable a particular log, specify the number of logs to be kept, and specify a format for log output.
 - Configuring Java logging using the administrative console

- If you do not want log and trace from Jakarta Commons Logging to use the WebSphere log and trace infrastructure, reconfigure the Jakarta Commons Logging.
 - “Configuring applications to use Jakarta Commons Logging” on page 33

Note: Use the WebSphere log and trace infrastructure for all of your log content to make problem source identification simpler.

- Restart the application server after making static configuration changes.

Example

The sample security policy that follows grants access to the file system and runtime classes. Include this security policy, with the entry permission `java.util.logging.LoggingPermission "control"`, in the META-INF directory of your application if you want your applications to programmatically alter controlled properties of loggers and handlers. The META-INF file is located in the following locations for the different module types:

Project name	Location
EJB projects	ejbModule/META-INF/MANIFEST.MF
Application client projects	appClientModule/META-INF/MANIFEST.MF
Dynamic web projects	WebContent/META-INF/MANIFEST.MF
Connector projects	connectorModule/META-INF/MANIFEST.MF

Below is a sample security policy that grants permission to modify logging properties:

```

////////////////////////////////////
//
// WebSphere Application Server Security Policy
//
////////////////////////////////////

////////////////////////////////////
// Allow all access to the file system and runtime classes
////////////////////////////////////
grant codeBase "file:${application}" {
    permission java.util.logging.LoggingPermission "control";
};

```

Using Java logging in an application

This topic describes how to use Java logging within an application.

About this task

To create an application using Java logging, perform the following steps:

Procedure

- Optional: Create the necessary handler, formatter, and filter classes if you need your own log files.

Note: Use the WebSphere log and trace infrastructure to make problem source identification simpler, rather than creating separate log files.

- Optional: If localized messages are used by the application, create a resource bundle, as described in “Creating log resource bundles and message files” on page 30.
- In the application code, get a reference to a logger instance, as described in “Using a logger” on page 19.

4. Insert the appropriate message and trace logging statements in the application, as described in “Using a logger.”

Using a logger

You can use Java logging to log messages and add tracing.

About this task

Java provides a log and trace package, `java.util.logging`, that you can use to instrument your applications. This topic provides recommendations about how to use the log and trace package.

Procedure

1. Use `WsLevel.DETAIL` level and above for messages, and lower levels for trace. The WebSphere Application Server Extension API (the `com.ibm.websphere.logging` package) contains the `WsLevel` class.

For messages use:

```
WsLevel.FATAL
Level.SEVERE
Level.WARNING
WsLevel.AUDIT
Level.INFO
Level.CONFIG
WsLevel.DETAIL
```

For trace use:

```
Level.FINE
Level.FINER
Level.FINEST
```

2. Use the `logp` method instead of the `log` or the `logrb` method. The `logp` method accepts parameters for class name and method name. The `log` and `logrb` methods will generally try to infer this information, but the performance penalty is prohibitive. In general, the `logp` method has less performance impact than the `log` or the `logrb` method.
3. Avoid using the `logrb` method. This method leads to inefficient caching of resource bundles and poor performance.
4. Use the `isLoggable` method to avoid creating data for a logging call that does not get logged. For example:

```
if (logger.isLoggable(Level.FINEST)) {
    String s = dumpComponentState(); // some expensive to compute method
    logger.logp(Level.FINEST, className, methodName, "componentX state
dump:\n{0}", s);
}
```

Localized messages

The following sample applies to localized messages:

```
// note - generally avoid use of FINE, FINER, FINEST levels for messages to be consistent with
// WebSphere Application Server

String componentName = "com.ibm.websphere.componentX";
String resourceBundleName = "com.ibm.websphere.componentX.Messages";
Logger logger = Logger.getLogger(componentName, resourceBundleName);

// "Convenience" methods - not generally recommended due to lack of class
// method names
// - cannot specify message substitution parameters
// - cannot specify class and method names
if (logger.isLoggable(Level.SEVERE))
    logger.severe("MSG_KEY_01");

if (logger.isLoggable(Level.WARNING))
    logger.warning("MSG_KEY_01");

if (logger.isLoggable(Level.INFO))
    logger.info("MSG_KEY_01");
```

```

if (logger.isLoggable(Level.CONFIG))
    logger.config("MSG_KEY_01");

// log methods are not generally used due to lack of class and method
names
// - enable use of WebSphere Application Server-specific levels
// - enable use of message substitution parameters
// - cannot specify class and method names
if (logger.isLoggable(WsLevel.FATAL))
    logger.log(WsLevel.FATAL, "MSG_KEY_01", "parameter 1");

if (logger.isLoggable(Level.SEVERE))
    logger.log(Level.SEVERE, "MSG_KEY_01", "parameter 1");

if (logger.isLoggable(Level.WARNING))
    logger.log(Level.WARNING, "MSG_KEY_01", "parameter 1");

if (logger.isLoggable(WsLevel.AUDIT))
    logger.log(WsLevel.AUDIT, "MSG_KEY_01", "parameter 1");

if (logger.isLoggable(Level.INFO))
    logger.log(Level.INFO, "MSG_KEY_01", "parameter 1");

if (logger.isLoggable(Level.CONFIG))
    logger.log(Level.CONFIG, "MSG_KEY_01", "parameter 1");

if (logger.isLoggable(WsLevel.DETAIL))
    logger.log(WsLevel.DETAIL, "MSG_KEY_01", "parameter 1");

// logp methods are the way to log
// - enable use of WebSphere Application Server-specific levels
// - enable use of message substitution parameters
// - enable use of class and method names
if (logger.isLoggable(WsLevel.FATAL))
    logger.logp(WsLevel.FATAL, className, methodName, "MSG_KEY_01",
"parameter 1");

if (logger.isLoggable(Level.SEVERE))
    logger.logp(Level.SEVERE, className, methodName, "MSG_KEY_01",
"parameter 1");

if (logger.isLoggable(Level.WARNING))
    logger.logp(Level.WARNING, className, methodName, "MSG_KEY_01",
"parameter 1");

if (logger.isLoggable(WsLevel.AUDIT))
    logger.logp(WsLevel.AUDIT, className, methodName, "MSG_KEY_01",
"parameter 1");

if (logger.isLoggable(Level.INFO))
    logger.logp(Level.INFO, className, methodName, "MSG_KEY_01",
"parameter 1");

if (logger.isLoggable(Level.CONFIG))
    logger.logp(Level.CONFIG, className, methodName, "MSG_KEY_01",
"parameter 1");

if (logger.isLoggable(WsLevel.DETAIL))
    logger.logp(WsLevel.DETAIL, className, methodName, "MSG_KEY_01",
"parameter 1");

// logrb methods are not generally used due to diminished performance
of switching resource bundles dynamically
// - enable use of WebSphere Application Server-specific levels
// - enable use of message substitution parameters
// - enable use of class and method names
String resourceBundleNameSpecial =
"com.ibm.websphere.componentX.MessagesSpecial";

if (logger.isLoggable(WsLevel.FATAL))
    logger.logrb(WsLevel.FATAL, className, methodName, resourceBundleNameSpecial,
"MSG_KEY_01", "parameter 1");

if (logger.isLoggable(Level.SEVERE))
    logger.logrb(Level.SEVERE, className, methodName, resourceBundleNameSpecial,
"MSG_KEY_01", "parameter 1");

if (logger.isLoggable(Level.WARNING))
    logger.logrb(Level.WARNING, className, methodName, resourceBundleNameSpecial,
"MSG_KEY_01", "parameter 1");

if (logger.isLoggable(WsLevel.AUDIT))
    logger.logrb(WsLevel.AUDIT, className, methodName, resourceBundleNameSpecial,
"MSG_KEY_01", "parameter 1");

if (logger.isLoggable(Level.INFO))
    logger.logrb(Level.INFO, className, methodName, resourceBundleNameSpecial,

```

```

"MSG_KEY_01", "parameter 1");

if (logger.isLoggable(Level.CONFIG))
    logger.logrb(Level.CONFIG, className, methodName, resourceBundleNameSpecial,
"MSG_KEY_01", "parameter 1");

if (logger.isLoggable(WebSocketLevel.DETAIL))
    logger.logrb(WebSocketLevel.DETAIL, className, methodName, resourceBundleNameSpecial,
"MSG_KEY_01", "parameter 1");

```

For trace, or content that is not localized, the following sample applies:

```

// note - generally avoid use of FATAL, SEVERE, WARNING, AUDIT,
// INFO, CONFIG, DETAIL levels for trace
// to be consistent with WebSphere Application Server

String componentName = "com.ibm.websphere.componentX";
Logger logger = Logger.getLogger(componentName);

// Entering / Exiting methods are used for non trivial methods
if (logger.isLoggable(Level.FINER))
    logger.entering(className, methodName);

if (logger.isLoggable(Level.FINER))
    logger.entering(className, methodName, "method param1");

if (logger.isLoggable(Level.FINER))
    logger.exiting(className, methodName);

if (logger.isLoggable(Level.FINER))
    logger.exiting(className, methodName, "method result");

// Throwing method is not generally used due to lack of message - use
// logp with a throwable parameter instead
if (logger.isLoggable(Level.FINER))
    logger.throwing(className, methodName, throwable);

// Convenience methods are not generally used due to lack of class
// method names
// - cannot specify message substitution parameters
// - cannot specify class and method names
if (logger.isLoggable(Level.FINE))
    logger.fine("This is my trace");

if (logger.isLoggable(Level.FINER))
    logger.finer("This is my trace");

if (logger.isLoggable(Level.FINEST))
    logger.finest("This is my trace");

// log methods are not generally used due to lack of class and
// method names
// - enable use of WebSphere Application Server-specific levels
// - enable use of message substitution parameters
// - cannot specify class and method names
if (logger.isLoggable(Level.FINE))
    logger.log(Level.FINE, "This is my trace", "parameter 1");

if (logger.isLoggable(Level.FINER))
    logger.log(Level.FINER, "This is my trace", "parameter 1");

if (logger.isLoggable(Level.FINEST))
    logger.log(Level.FINEST, "This is my trace", "parameter 1");

// logp methods are the recommended way to log
// - enable use of WebSphere Application Server-specific levels
// - enable use of message substitution parameters
// - enable use of class and method names
if (logger.isLoggable(Level.FINE))
    logger.logp(Level.FINE, className, methodName, "This is my trace",
"parameter 1");

if (logger.isLoggable(Level.FINER))
    logger.logp(Level.FINER, className, methodName, "This is my trace",
"parameter 1");

if (logger.isLoggable(Level.FINEST))
    logger.logp(Level.FINEST, className, methodName, "This is my trace",
"parameter 1");

// logrb methods are not applicable for trace logging because no localization
// is involved

```

There may be occasions when you want to propagate log records to your own log handlers rather than participate in integrated logging. To use a stand-alone log handler, set the `useParentHandlers` flag to `false` in your application. The mechanism for creating a custom handler is the Handler class support that is provided by the IBM Developer Kit, Java Technology Edition. If you are not familiar with handlers, as implemented by the Developer Kit, you can get more information from various texts, or by reading the API documentation for the `java.util.logging` API. The following sample shows a custom handler:

```
import java.io.FileOutputStream;
import java.io.PrintWriter;
import java.util.logging.Handler;
import java.util.logging.LogRecord;

/**
 * MyCustomHandler outputs contents to a specified file
 */
public class MyCustomHandler extends Handler {

    FileOutputStream fileOutputStream;
    PrintWriter printWriter;

    public MyCustomHandler(String filename) {
        super();

        // check input parameter
        if (filename == null || filename == "")
            filename = "mylogfile.txt";

        try {
            // initialize the file
            fileOutputStream = new FileOutputStream(filename);
            printWriter = new PrintWriter(fileOutputStream);
            setFormatter(new SimpleFormatter());
        }
        catch (Exception e) {
            // implement exception handling...
        }
    }

    /* (non-API documentation)
     * @see java.util.logging.Handler#publish(java.util.logging.LogRecord)
     */
    public void publish(LogRecord record) {
        // ensure that this log record should be logged by this Handler
        if (!isLoggable(record))
            return;

        // Output the formatted data to the file
        printWriter.println(getFormatter().format(record));
    }

    /* (non-API documentation)
     * @see java.util.logging.Handler#flush()
     */
    public void flush() {
        printWriter.flush();
    }

    /* (non-API documentation)
     * @see java.util.logging.Handler#close()
     */
    public void close() throws SecurityException {
        printWriter.close();
    }
}
```


A custom filter provides optional, secondary control over what is logged, beyond the control that is provided by the level. The mechanism for creating a custom filter is the Filter interface support that is provided by the IBM Developer Kit, Java Technology Edition. If you are not familiar with filters, as implemented by the Developer Kit, you can get more information from various texts, or by reading the API documentation for the java.util.logging API.

The following example shows a custom filter:

```
/**
 * This class filters out all log messages starting with SECJ022E, SECJ0373E, or SECJ0350E.
 */
import java.util.logging.Filter;
import java.util.logging.Handler;
import java.util.logging.Logger;
import java.util.logging.LogRecord;

public class MyFilter implements Filter {
    public boolean isLoggable(LogRecord lr) {
        String msg = lr.getMessage();
        if (msg.startsWith("SECJ022E") || msg.startsWith("SECJ0373E") || msg.startsWith("SECJ0350E")) {
            return false;
        }
        return true;
    }
}

//This code will register the above log filter with the root Logger's handlers (including the WAS system logs):
...
Logger rootLogger = Logger.getLogger("");
rootLogger.setFilter(new MyFilter());
```

A formatter formats events. Handlers are associated with one or more formatters. The mechanism for creating a custom formatter is the Formatter class support that is provided by the IBM Developer Kit, Java Technology Edition. If you are not familiar with formatters, as implemented by the Developer Kit, you can get more information from various texts, or by reading the API documentation for the java.util.logging API.

The following example shows a custom formatter:

```
import java.util.Date;
import java.util.logging.Formatter;
import java.util.logging.LogRecord;

/**
 * MyCustomFormatter formats the LogRecord as follows:
 * date level localized message with parameters
 */
public class MyCustomFormatter extends Formatter {

    public MyCustomFormatter() {
        super();
    }

    public String format(LogRecord record) {

        // Create a StringBuffer to contain the formatted record
        // start with the date.
        StringBuffer sb = new StringBuffer();

        // Get the date from the LogRecord and add it to the buffer
        Date date = new Date(record.getMillis());
        sb.append(date.toString());
        sb.append(" ");

        // Get the level name and add it to the buffer
        sb.append(record.getLevel().getName());
        sb.append(" ");

        // Get the formatted message (includes localization
        // and substitution of parameters) and add it to the buffer
        sb.append(formatMessage(record));
        sb.append("\n");
    }
}
```

```

    return sb.toString();
}
}

```

Adding custom handlers, filters, and formatters enables you to customize your logging environment beyond what can be achieved by the configuration of the default WebSphere Application Server logging infrastructure. The following example demonstrates how to add a new handler to process requests to the `com.myCompany` subtree of loggers (see “Configuring the logger hierarchy” on page 29). The main method in this sample gives an example of how to use the newly configured logger.

```

import java.util.Vector;
import java.util.logging.Filter;
import java.util.logging.Formatter;
import java.util.logging.Handler;
import java.util.logging.Level;
import java.util.logging.Logger;

public class MyCustomLogging {

    public MyCustomLogging() {
        super();
    }

    public static void initializeLogging() {

        // Get the logger that you want to attach a custom Handler to
        String defaultResourceBundleName = "com.myCompany.Messages";
        Logger logger = Logger.getLogger("com.myCompany", defaultResourceBundleName);

        // Set up a custom Handler (see MyCustomHandler example)
        Handler handler = new MyCustomHandler("MyOutputFile.log");

        // Set up a custom Filter (see MyCustomFilter example)
        Vector acceptableLevels = new Vector();
        acceptableLevels.add(Level.INFO);
        acceptableLevels.add(Level.SEVERE);
        Filter filter = new MyCustomFilter(acceptableLevels);

        // Set up a custom Formatter (see MyCustomFormatter example)
        Formatter formatter = new MyCustomFormatter();

        // Connect the filter and formatter to the handler
        handler.setFilter(filter);
        handler.setFormatter(formatter);

        // Connect the handler to the logger
        logger.addHandler(handler);

        // avoid sending events logged to com.myCompany showing up in WebSphere
        // Application Server logs
        logger.setUseParentHandlers(false);
    }

    public static void main(String[] args) {
        initializeLogging();

        Logger logger = Logger.getLogger("com.myCompany");

        logger.info("This is a test INFO message");
        logger.warning("This is a test WARNING message");
        logger.logp(Level.SEVERE, "MyCustomLogging", "main", "This is a test SEVERE message");
    }
}

```

When the above program is run, the output of the program is written to the `MyOutputFile.log` file. The content of the log is in the expected log file, as controlled by the custom handler, and is formatted as defined by the custom formatter. The warning message is filtered out, as specified by the configuration of the custom filter. The output is as follows:

```
C:\>type MyOutputFile.log
Sat Sep 04 11:21:19 EDT 2004 INFO This is a test INFO message
Sat Sep 04 11:21:19 EDT 2004 SEVERE This is a test SEVERE message
```

Loggers

Loggers are used by applications and runtime components to capture message and trace events.

When situations occur that are significant either due to a change in state, for example when a server completes startup or because a potential problem is detected, such as a timeout waiting for a resource, a message is written to the logs. Trace events are logged in debugging scenarios, where a developer needs a clear view of what is occurring in each component to understand what might be going wrong. Logged events are often the only events available when a problem is first detected, and are used during both problem recovery and problem resolution.

Loggers are organized hierarchically. Each logger can have zero or more child loggers.

Loggers can be associated with a resource bundle. If specified, the resource bundle is used by the logger to localize messages that are logged to the logger. If the resource bundle is not specified, a logger uses the same resource bundle as its parent.

You can configure loggers with a level. If specified, the level is compared by the logger to incoming events. The events that are less severe than the level set for the logger are ignored by the logger. If the level is not specified, a logger takes on the level that is used by its parent. The default level for loggers is `Level.INFO`.

Loggers can have zero or more attached handlers. If supplied, all events that are logged to the logger are passed to the attached handlers. Handlers write events to output destinations such as log files or network sockets. When a logger finishes passing a logged event to all of the handlers that are attached to that logger, the logger passes the event to the handlers that are attached to the parents of the logger. This process stops if a parent logger is configured not to use its parent handlers. Handlers in WebSphere Application Server are attached to the root logger. Set the `useParentHandlers` logger property to `false` to prevent the logger from writing events to handlers that are higher in the hierarchy.

Loggers can have a filter. If supplied, the filter is invoked for each incoming event to tell the logger whether or not to ignore it.

Applications interact directly with loggers to log events. To obtain or create a logger, a call is made to the `Logger.getLogger` method with a name for the logger. Typically, the logger name is either the package qualified class name or the name of the package that the logger is used by. The hierarchical logger namespace is automatically created by using the dots in the logger name. For example, the `com.ibm.websphere.ras` logger has a `com.ibm.websphere` parent logger, which has a `com.ibm` parent. The parent at the top of the hierarchy is referred to as the *root logger*. This root logger is created during initialization. The root logger is the parent of the `com` logger.

Loggers are structured in a hierarchy. Every logger, except the root logger, has one parent. Each logger can also have 0 or more children. A logger inherits log handlers, resource bundle names, and event filtering settings from its parent in the hierarchy. The logger hierarchy is managed by the `LogManager` function.

Loggers create log records. A log record is the container object for the data of an event. This object is used by filters, handlers, and formatters in the logging infrastructure.

The logger provides several sets of methods for generating log messages. Some log methods take only a level and enough information to construct a message. Other, more complex logp (log precise) methods support the caller in passing class name and method name attributes, in addition to the level and message information. The logrb (log with resource bundle) methods add the capability of specifying a resource bundle as well as the level, message information, class name, and method name. Using methods such as severe, warning, fine, finer, and finest you can log a message at a particular level. For more information on logging and how to use it in your applications read “Using Java logging in an application” on page 18. For a complete list of methods, see the java.util.logging documentation at <http://java.sun.com/javase/>.

Log handlers

Log handlers write log record objects to output devices like log files, sockets, and notification mechanisms.

Loggers can have zero or more attached handlers. All objects that are logged to the logger are passed to the attached handlers, if handlers are supplied.

You can configure handlers with a level. The handler compares the level that is specified in the logged object to the level that is specified for the handler. If the level of the logged object is less severe than the level set in the handler, the object is ignored by the handler. The default level for handlers is ALL.

Handlers can have a filter. If a filter is supplied, the filter is invoked for each incoming object to tell the handler whether or not to ignore it.

Handlers can have a formatter. If a formatter is supplied, the formatter controls how the logged objects are formatted. For example, the formatter can decide to first include the time stamp, followed by a string representation of the level, followed by the message that is included in the logged object. The handler writes this formatted representation to the output device.

Both loggers and handlers can have levels and filters, and a logged object must pass all of these elements to be output. For example, you can set the logger level to FINE, but if the handler level is set at WARNING, only WARNING level messages are displayed in the output for that handler. Conversely, if your log handler is set to output all messages (level=All), but the logger level is set to WARNING, the logger never sends messages beneath the WARNING to the log handler.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log , SystemErr.log, trace.log, and activity.log files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Log levels

Levels control which events are processed by Java logging. WebSphere Application Server controls the levels of all loggers in the system.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log , SystemErr.log, trace.log, and activity.log files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

The level value is set from configuration data when the logger is created and can be changed at run time from the administrative console. If a level is not set in the configuration data, a level is obtained by proceeding up the hierarchy until a parent with a level value is found. You can also set a level for each

handler to indicate which events are published to an output device. When you change the level for a logger in the administrative console, the change is propagated to the children of the logger.

Levels are cumulative; a logger can process logged objects at the level that is set for the logger, and at all levels above the set level.

Table 3. Valid log levels. This table lists valid logging levels.

Level	Content / Significance
Off	No events are logged.
Fatal	Task cannot continue and component cannot function.
Severe	Task cannot continue, but component can still function
Warning	Potential error or impending error
Audit	Significant event affecting server state or resources
Info	General information outlining overall task progress
Config	Configuration change or status
Detail	General information detailing subtask progress
Fine	Trace information - General trace
Finer	Trace information - Detailed trace + method entry / exit / return values
Finest	Trace information - A more detailed trace - Includes all the detail that is needed to debug problems
All	All events are logged. If you create custom levels, All includes your custom levels, and can provide a more detailed trace than Finest.

For instructions on how to set logging levels, read the topic about configuring Java logging using the administrative console.

Note: Trace information, which includes events at the Fine, Finer and Finest levels, can be written only to the trace log. Therefore, if you do not enable diagnostic trace, setting the log detail level to Fine, Finer, or Finest does not effect the logged data.

Log filters

Log filters help control more detailed logging settings that are not handled by usual log level settings.

A filter provides an optional, secondary control over what is logged, beyond the control that is provided by setting the level. Applications can apply a filter mechanism to control logging output through the logging APIs. An example of filter usage is to suppress all the events with a particular message key.

A filter is attached to a logger or log handler using the appropriate `setFilter` method. For a complete list of filter methods, see the `java.util.logging` documentation at <http://java.sun.com/javase/>

Log formatters

Log formatters format log messages so they can be used by various log handlers.

Handlers can be configured with a log formatter that knows how to format log records. The event, which is represented by the log record object, is passed to the appropriate formatter by the handler. The formatter returns formatted output to the handler, which writes the output to the output device.

The formatter is responsible for rendering the event for output. This formatter uses the resource bundle that is specified in the event to look up the message in the appropriate language.

Formatters are attached to handlers using the `setFormatter` method.

IBM i You can select a formatter for a handler using the administrative console panels. See Diagnostic trace service settings for details.

You can find the `java.util.logging` documentation at <http://java.sun.com/javase/>.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Java logging

Java logging is the logging toolkit that is provided by the `java.util.logging` package. Java logging provides a standard logging API for your applications.

Message logging (messages) and diagnostic trace (trace) are conceptually similar, but do have important differences. These differences are important for application developers to understand to use these tools properly. The following operational definitions of messages and trace are provided.

Message

A message entry is an informational record that is intended for end users, systems administrators, and support personnel to view. The text of the message must be clear, concise, and interpretable by an end user. Messages are typically localized and displayed in the national language of the end user. Although the destination and lifetime of messages might be configurable, enable some level of message logging in normal system operation. Use message logging judiciously because of performance considerations and the size of the message repository.

Trace A trace entry is an information record that is intended for service engineers or developers to use. As such, a trace record might be considerably more complex, verbose, and detailed than a message entry. Localization support is typically not used for trace entries. Trace entries can 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 can be enabled as needed to gather diagnostic information.

The application server redirects the system streams at the server startup. There is no way to allow the application to output logging to the console because the system streams can not be obtained by the application. If you would like to use console to monitor the application without using the console handler, you can either monitor the `SystemOut.log` file, or monitor a file created by another file handler.

Note: The application server uses Java logging internally and therefore certain restrictions apply for using system streams with this logging API by applications. During server startup, the standard output and error streams are replaced with special streams that write to the logging infrastructure, in order to include the output of the system streams in the log files. Because of this, applications can not use `java.util.logging.ConsoleHandler`, or any handler writing to `SystemErr.log` or `System.out` streams, attached to the root logger. If the user does attach the handler to the root logger, an infinite loop is created within the logging infrastructure, leading to stack overflow and server crash.

If the use of a handler that writes to system streams is necessary, attach it to a non-root logger so that it does not publish log records to parent handlers. The data written to the system streams is then formatted and written to the corresponding system stream log file. To monitor what is being written system streams, the configured log files (`SystemOut.log` and `SystemErr.log` by default) can be monitored.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace

infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Configuring the logger hierarchy

WebSphere Application Server handlers are attached to the Java root logger, which is at the top of the logger hierarchy. As a result, any request from anywhere in the logger tree can be processed by WebSphere Application Server handlers.

About this task


You can configure your application server to handle logs in many different ways. Configure your log settings based upon your configuration and the logging structure that best suits your needs.

Procedure

- Forward all application logging requests to the WebSphere Application Server handlers. This behavior is the default.
- Forward all application logging requests to your own custom handlers. Set the `useParentHandlers` option to `false` on one of your custom loggers, and then attach your handlers to that logger.
- Forward all application logging requests to both WebSphere Application Server handlers, and your custom handlers, but do not forward WebSphere Application Server logging requests to your custom handlers. Set the `useParentHandlers` option to `true` on one of your non-root custom loggers, and then attach your handlers to that logger. `True` is the default setting.
- Forward all WebSphere Application Server logging requests to both WebSphere Application Server handlers, and your custom handlers. Logging requests are always forwarded to WebSphere Application Server handlers. To forward WebSphere Application Server requests to your custom handlers, attach your custom handlers to the Java root logger, so that they are at the same level in the hierarchy as the WebSphere Application Server handlers.

Example

The following example shows how these requirements can be met using the Java logging infrastructure:



Creating log resource bundles and message files

You can forward messages that are written to the internal WebSphere Application Server logs to other processes for display. Messages that are displayed on the administrative 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.

About this task

Every method that accepts messages localizes those messages. The mechanism for providing localized messages is the resource bundle support provided by the IBM Developer Kit, Java Technology Edition. If you are not familiar with resource bundles as implemented by the Developer Kit, you can get more information from various texts, or by reading the API documentation for the `java.util.ResourceBundle`, `java.util.ListResourceBundle` and `java.util.PropertyResourceBundle` classes, as well as the `java.text.MessageFormat` class.

The `PropertyResourceBundle` class is the preferred mechanism to use.

To properly localize the message, the displaying process must have access to the resource bundle where the message text is stored. You must package the resource bundle separately from the application, and install it in a location where the viewing process can access it.

By default, the WebSphere Application Server runtime localizes all the messages when they are logged. This localization eliminates the need to pass a `.jar` file to the application, unless you need to localize in a different location. However, you can use the early binding technique to localize messages as they log. An application that uses early binding must localize the message before logging it. The application looks up

the localized text in the resource bundle and formats the message. Use the early binding technique to package the application resource bundles with the application.

To create a resource bundle, perform the following steps.

Procedure

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 white space only, 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 of white space only, 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 process is not recommended, because escaping characters is error prone and confusing. Instead, use a valid separator character that does not display 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 the construction of properties files.

2. Translate the file 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, put the bundle in a directory that is part of the application class path.
4. When a message logger is obtained from the log manager, configure it to use a particular resource bundle. Messages logged with the Logger API use this resource bundle when message localization is performed. At run time, the user locale setting determines the properties file from which to extract the message that is specified by a message key, 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.

Example

You can create resource bundles in several ways. The best and easiest way is to create a properties file that supports a properties resource bundle. This example 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 are inserted. All the normal properties file conventions and rules apply to this file. In addition, the creator must be aware of other restrictions that are imposed on the values by the Java `MessageFormat` class. For example, apostrophes must be escaped or they cause a problem. Avoid the use of non-portable characters. WebSphere Application Server does not support the use 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 is in the class path. Assume that the properties file is stored in the subdirectory `baseDir` that is not in the class path (for example, `baseDir/subDir1/subDir2/resources`). To allow the messages file to resolve, the `subDir1.subDir2.resources.DefaultMessage` name is used to identify the property resource bundle and is passed to the message logger.

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

```
# Contents of the 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}
```

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

What to do next

The application locates the resource bundle based on the file location relative to any directory in the class path. For instance, if the `DefaultMessages.properties` property resource bundle 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.

Logger.properties file for configuring logger settings

Use the `Logger.properties` file to set logger attributes for specific loggers.

The properties file is loaded the first time that the `Logger.getLogger(logger_name)` method is called within an application.

Important: The name of the `Logger.properties` file is case sensitive. Use a capital "L" in the file name.

When an application calls the `Logger.getLogger` method for the first time, all the available logger properties files are loaded. Applications can provide `Logger.properties` files in:

- the META-INF directory of the Java archive (JAR) file for the application
- directories included in the class path of an application module
- directories included in the application class path

The properties file contains two categories of parameters, logger control and logger data:

- Logger control information
 - Minimum localization level: The minimum `LogRecord` level for which localization is attempted
 - Group: The logical group that this component belongs to
 - Event factory: The Common Base Event template file to use with the event factory. The naming convention for this template is the fully qualified component name, with a file extension of `.event.xml`. For example, a template that applies to the `com.ibm.compXYZ` package is called `com.ibm.compXYZ.event.xml`.
- Logger data information
 - Product name
 - Organization name
 - Component name
 - Extensions and additional properties

Syntax of the Logger.properties file

Use the following syntax to set logger properties:

```
<logger base name>.<property>=value
```

where:

logger base name is the starting part of the logger name to which the property applies. All loggers with names starting with this string have the property applied.

property is one of the following properties:

- organization
- product
- component
- minimum_localization_level
- group
- eventfactory

Sample Logger.properties file

In the following sample, the `com.ibm.xyz.MyEventFactory` event factory is used by any loggers in the `com.ibm.websphere.abc` package or any sub packages that do not override this value in their configuration file.

```
com.ibm.websphere.abc.eventfactory=com.ibm.xyz.MyEventFactory
```

Group Logger.properties file

In the following example, the group is `MyTraceGroup` and the components are `com.ibm.stuff` and `com.ibm.morestuff`:

```
com.ibm.stuff.group=MyTraceGroup
com.ibm.morestuff.group=MyTraceGroup
```

Configuring applications to use Jakarta Commons Logging

Jakarta Commons Logging provides a simple logging interface and thin wrappers for several logging systems. WebSphere Application Server supports Jakarta Commons Logging by providing a logger. The support does not change interfaces defined by Jakarta Commons Logging.

Before you begin

The WebSphere Application Server logger is a thin wrapper for the WebSphere Application Server logging facility. The logger name is `com.ibm.websphere.commons.logging.WsJDK14Logger`. The logger can handle logging objects defined by either of the following:

- Java Logging found in Java Specification Request 47: Logging API Specification
- Common Base Event

A *logging object* is an object that holds logging entry information.

To better understand Jakarta Commons Logging, read Jakarta Commons and the specifications for Java Logging and for Common Base Event. To better understand use of the WebSphere Application Server logger, read “Jakarta Commons Logging” on page 34.

About this task

WebSphere Application Server provides the Jakarta Commons Logging binary distribution in its `libraries` directory. By default, the product uses the Jakarta Commons Logging `LogFactory` implementation and `JDK14Logger`.

best-practices: The default configuration of Jakarta Commons Logging is stored in the `commons-logging.properties` file. To specify the factory class to use with Jakarta Commons Logging in an application, provide a file named

`org.apache.commons.logging.LogFactory`, located in `META-INF/services` directory, that contains the name of the factory class on the first line. This is the configuration mechanism for the JAR file service provider, as defined in JDK 1.3 and above.

For an application to use the WebSphere Application Server logger, the application must provide its own configuration for the logger. To configure an application to use the WebSphere Application Server logger, complete the steps that follow.

Procedure

1. Examine “Configurations for the WebSphere Application Server logger” on page 37 and determine which configuration best suits your application.
2. Change your application configuration as needed to enable use of the WebSphere Application Server logger.

Results

After the application starts, Jakarta Commons Logging routes the application's logging output to the WebSphere Application Server logger.

Jakarta Commons Logging

Jakarta Commons Logging provides a simple logging interface and thin wrappers for several logging systems. The logging interface enables application logging to be simple and independent of the logging system that the application uses. You can change the logging implementation for a deployed application without having to change the application logging code. However, the simplicity of the logging interface prevents the application from leveraging all the functionality of the logging systems.

This topic provides the following information about Jakarta Commons Logging in WebSphere Application Server:

- “Support for Jakarta Commons Logging”
- “Benefits of support for Jakarta Commons Logging”
- “Overview of the process for using Jakarta Commons Logging” on page 35
- “Classes used to obtain a logger factory and logger” on page 35
- “Logger level configuration and mapping” on page 36

Support for Jakarta Commons Logging

The product supports Jakarta Commons Logging by providing a logger, a thin wrapper for the WebSphere Application Server logging facility. The logger can handle both Java Logging (JSR-47) and Common Base Event logging objects. A *logging object* is an object that holds logging entry information.

The product support for Jakarta Commons Logging does not change interfaces defined by Jakarta Commons Logging.

Benefits of support for Jakarta Commons Logging

The WebSphere Application Server support for Jakarta Commons Logging provides the following benefits:

- WebSphere Application Server is pre-configured to use Jakarta Commons Logging.
All of the functionality of Jakarta Commons Logging is provided for any application or WebSphere Application Server component. Logging calls are routed by default to the underlying WebSphere Application Server logging facility.
- A logger that uses the WebSphere Application Server logging facility.

Applications and components can pass both Java Logging and Common Base Event logging objects to the WebSphere Application Server logger without conversion to strings, providing applications with enhanced logging. Further, Jakarta Commons Logging Logger levels are integrated into WebSphere Application Server administrative facilities.

Overview of the process for using Jakarta Commons Logging

Logging with Jakarta Commons Logging consists of the steps that follow. “Configurations for the WebSphere Application Server logger” on page 37 provides details on configuring your application to use the WebSphere Application Server logger.

1. Obtain an instance of a logger factory.

To obtain a logger factory, use Jakarta Commons Logging code. You can configure the code to meet your needs. In WebSphere Application Server, Jakarta Commons Logging is configured by default to instantiate the Jakarta Commons Logging default logger factory. Applications or WebSphere Application Server components can provide their own configuration if they use a different logger factory implementation. Applications can use more than one factory.

2. Obtain an instance of a logger.

To obtain a logger, use code implemented by a logger factory. Configuration of the code is implementation specific.

The WebSphere Application Server logger implements the methods defined in the logging interface. The logging methods take at least one argument, which can be any Java object. The WebSphere Application Server logger, the `WsJDK14Logger` logger described in “Classes used to obtain a logger factory and logger,” handles the following objects passed into the following logging methods:

CommonBaseEvent

Wrapped into `CommonBaseEventLogRecord`

CommonBaseEventLogRecord

Passed without change

LogRecord

Passed without change

Other objects

Converted to String

Applications or WebSphere Application Server components can provide their own configuration if they use an implementation of a logger that is not specific to WebSphere Application Server. An application must know what factory is being used in order to configure it.

3. Start your application. Jakarta Commons Logging routes the application's logging output to the designated logger

Classes used to obtain a logger factory and logger

Table 4. Jakarta Commons Logging class descriptions. Use the classes for a logger factory instance and logger.

Class name	Description
LogFactory	<p><i>LogFactory</i> is a Jakarta Commons Logging class that implements initialization logic. <i>LogFactory</i> is an abstract class that every logger factory implementation has to extend. It provides static methods for obtaining:</p> <ul style="list-style-type: none"> • An instance of a factory class • Instances of a logger, using an instance of the factory class <p><i>LogFactory</i> provides methods for obtaining instances of loggers, although these methods delegate the logger instantiation and configuration to an instance of a logger factory class.</p> <p>Logger factories, once instantiated, are cached on a per context class loader basis. The instances in a cache can be released. This functionality is designed for platform container implementations rather than for applications.</p>

Table 4. Jakarta Commons Logging class descriptions (continued). Use the classes for a logger factory instance and logger.

Class name	Description
LogFactoryImpl	<i>LogFactoryImpl</i> is a Jakarta Commons Logging concrete class that implements the default logger factory using methods in <i>LogFactory</i> . To use Java Logging, there must always be at least one instance of a logger factory class, even if the application has not explicitly obtained one. If the configuration does not name a logger factory class, <i>LogFactoryImpl</i> is used as the default.
Log	<i>Log</i> is a Jakarta Commons Logging interface for loggers. Commons logging loggers have to implement the <i>Log</i> interface. Because the goal of Jakarta Commons Logging is to wrapper any logging system, the <i>Log</i> interface defines a small set of common logging methods. In WebSphere Application Server, <i>WsJDK14Logger</i> implements the <i>Log</i> interface. Logger instantiation and configuration is specific to every logger factory. Logging in WebSphere Application Server uses the default logger factory provided in Jakarta Commons Logging, which keeps instantiated loggers in cache, on a per context class loader basis.
WsJDK14Logger	<i>WsJDK14Logger</i> is a WebSphere Application Server class that provides a Jakarta Commons Logging logger by implementing the <i>Log</i> interface. The <i>WsJDK14Logger</i> logger differs from the Java Logging logger in that the <i>WsJDK14Logger</i> logger enables Java Logging or Common Base Event objects to be passed over without converting them into String objects. This prevents any information loss the conversion to String might cause as well as allows the logging output to be more descriptive and precise. In contrast, the Java Logginglogger that is provided in Jakarta Commons Logging converts objects passed into the logging calls to String objects prior to passing them over to the underlying Java Logging.

Logger level configuration and mapping

Because Jakarta Commons Logging loggers are thin wrappers for specific logging systems, the loggers do not have their own level, but use the level of the logger from the underlying logging system. Although the underlying system can provide methods for changing level, there are no methods for changing level defined on the *Log* interface, which all Jakarta Commons Logging loggers must implement. *WsJDK14Logger* uses the level of its underlying Java Logging logger.

Following table shows, on the left, the mapping of Jakarta Commons Logging levels within *WsJDK14Logger* to levels in the WebSphere Application Server implementation of Java Logging. The first column shows the levels defined in Java Logging and the level mapping in the Jakarta Commons Logging *JDK14Logger* to the Java Logging levels.

Table 5. Mapping of *WsJDK14Logger* levels to Java Logging levels. Compare the logging levels.

WsJDK14Logger	Java Logging in WebSphere Application Server	Java Logging	JDK14Logger
Fatal	Fatal		
Error	Severe	Severe	Fatal, Error
Warning	Warning	Warning	Warning
	Audit		
Info	Info	Info	Info
	Config	Config	
	Detail		
Debug	Fine	Fine	Debug
	Finer	Finer	
Trace	Finest	Finest	Trace

The WsJDK14Logger level is synchronized with the underlying Java Logging logger level. WebSphere Application Server administration controls the WsJDK14Logger level.

Configurations for the WebSphere Application Server logger

This topic describes several ways to configure an application to use the WebSphere Application Server logger.

The type of configuration that best suits an application depends upon the following:

- Whether the class loader order setting for the application is `Classes loaded with parent class loader first (Parent First)` or `Classes loaded with application class loader first (Parent Last)`, you can set the class loader delegation mode on a console page. For more details about class load order and delegation, consult the class loading chapter in the *Developing applications* PDF book.
- Whether Jakarta Commons Logging is bundled with the application configuration
- Whether Jakarta Commons Logging is provided within the application

The following tables describe the conditions required to enable an application to use the WebSphere Application Server logger.

Class loader mode is Parent First and Jakarta Commons Logging is bundled with the application

Table 6. Conditions required to use logger. When Parent First and Jakarta Commons Logging is bundled with an application.

Jakarta Commons Logging configuration	LogFactory instance	Log instance	Comments
<p>The application provides the configuration by either of the following:</p> <ul style="list-style-type: none"> • The properties file <code>commons-logging.properties</code> in the application classpath is not read by the LogFactory because the parent class loader finds the WebSphere properties file first. • The class name is read from the file <code>META-INF/services/org.apache.commons.logging.LogFactory</code> 	<p>The log factory used is the LogFactory implementation specified in the WebSphere Application Server default configuration, unless the configuration is provided in a META-INF file of the application or module.</p>	<p>The log used is either of the following:</p> <ul style="list-style-type: none"> • The Log implementation specified in the WebSphere Application Server default configuration • An application-specific Log implementation if an application-specific LogFactory that instantiates a different Log implementation is used. 	<p>The application parent class loader is the first class loader to load the Jakarta Commons Logging code. The WebSphere bundle that supports Jakarta Commons Logging provides the LogFactory static code that looks up the LogFactory configuration attributes.</p> <p>For the static LogFactory code to instantiate the LogFactory instance specified in the application configuration, the LogFactory instance must be on the classpath of the parent class loader.</p>
<p>Not provided by the application</p>	<p>The log factory used is the LogFactory implementation specified in the WebSphere default configuration.</p>	<p>The log used is the Log implementation specified in the WebSphere default configuration.</p>	<p>The Jakarta Commons Logging bundled with the application is not used.</p>

Class loader mode is Parent First and Jakarta Commons Logging is not bundled with the application

Table 7. Conditions required to use logger. When Parent First and Jakarta Commons Logging is not bundled with an application.

Jakarta Commons Logging configuration	LogFactory instance	Log instance	Comments
<p>The application provides the configuration by either of the following:</p> <ul style="list-style-type: none"> The properties file <code>commons-logging.properties</code> in the application classpath is not read by the LogFactory because the parent class loader finds the WebSphere Application Server properties file first. The class name is read from the file <code>META-INF/services/org.apache.commons.logging.LogFactory</code> 	<p>The log factory used is the LogFactory implementation specified in the WebSphere Application Server default configuration, unless the configuration is provided in a META-INF file of the application or module.</p>	<p>The log used is either of the following:</p> <ul style="list-style-type: none"> The Log implementation specified in the WebSphere Application Server default configuration An application-specific Log implementation if an application-specific LogFactory that instantiates a different Log implementation is used. 	<p>The application parent class loader is the first class loader to load the Jakarta Commons Logging code. The WebSphere bundle that supports Jakarta Commons Logging provides the LogFactory static code that looks up the LogFactory configuration attributes.</p> <p>For the static LogFactory code to instantiate the LogFactory instance specified in the application configuration, the LogFactory instance must be on the classpath of the parent class loader.</p>
<p>Not provided by the application</p>	<p>The log factory used is the LogFactory implementation specified in the WebSphere Application Server default configuration.</p>	<p>The log used is the Log implementation specified in the WebSphere Application Server default configuration.</p>	<p>Same as in the previous row</p>

Class loader mode is Parent Last and Jakarta Commons Logging is bundled with the application

Table 8. Conditions required to use logger. When Parent Last and Jakarta Commons Logging is bundled with an application.

Jakarta Commons Logging configuration	LogFactory instance	Log instance	Comments
<p>The application provides the configuration by either of the following:</p> <ul style="list-style-type: none"> The properties file <code>commons-logging.properties</code> in the application classpath is read by the LogFactory because the class loader finds the application properties file first. The class name is read from the file <code>META-INF/services/org.apache.commons.logging.LogFactory</code> 	<p>The log factory used is either of the following:</p> <ul style="list-style-type: none"> The default Jakarta Commons Logging LogFactory The LogFactory specified in the application configuration 	<p>The log used is the Log implementation specified in the application configuration.</p> <p>If the log factory used is the default Jakarta Commons Logging LogFactory, the Log implementation must be on the classpath of the application class loader.</p>	<p>The application class loader is the first class loader to load the Jakarta Commons Logging code. The application bundle that supports Jakarta Commons Logging provides the LogFactory static code that looks up the LogFactory configuration attributes.</p> <p>For the static LogFactory code to instantiate the LogFactory instance specified in the application configuration, the LogFactory instance must be on the classpath of the application class loader.</p>
<p>Not provided by the application</p>	<p>The log factory used is the LogFactory implementation specified in the WebSphere Application Server default configuration.</p>	<p>The log used is the Log implementation specified in the WebSphere Application Server default configuration.</p>	

Class loader mode is Parent Last and Jakarta Commons Logging is not bundled with the application

Table 9. Conditions required to use logger. When Parent Last and Jakarta Commons Logging is not bundled with an application.

Jakarta Commons Logging configuration	LogFactory instance	Log instance	Comments
<p>The application provides the configuration by either of the following:</p> <ul style="list-style-type: none"> The properties file <code>commons-logging.properties</code> in the application classpath is read by the LogFactory because the class loader finds the application properties file first. The class name is read from the file <code>META-INF/services/org.apache.commons.logging.LogFactory</code> 	<p>The log factory used is either of the following:</p> <ul style="list-style-type: none"> The default Jakarta Commons Logging LogFactory The LogFactory specified in the application configuration 	<p>The log used is the Log implementation specified in the application configuration.</p> <p>If the log factory used is the default Jakarta Commons Logging LogFactory, the Log implementation must be on the classpath of the application class loader.</p>	<p>There is no Jakarta Commons Logging code at the application class loader. Thus, the WebSphere bundle that supports Jakarta Commons Logging provides the LogFactory static code that looks up the LogFactory configuration attributes.</p> <p>For the static LogFactory code to instantiate the LogFactory instance specified in the application configuration, the LogFactory instance must be on the classpath of the parent class loader.</p>
Not provided by the application	The log factory used is the LogFactory implementation specified in the WebSphere Application Server default configuration.	The log used is the Log implementation specified in the WebSphere Application Server default configuration.	

Programming with the JRas framework

Use the JRas extensions to incorporate message logging and diagnostic trace into WebSphere Application Server applications.

Before you begin

Note: The JRas framework that is described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

About this task

The JRas extensions allow message logging and diagnostic trace to work with WebSphere Application Server applications. They are based on the stand-alone JRas logging toolkit.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Procedure

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.

JRas logging toolkit

The JRas logging toolkit provides diagnostic information to help the administrator diagnose problems or tune application performance.

Note: The JRes framework that is described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

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 failed and provide information. This action enables the administrator to take the proper corrective action. Those who develop or maintain applications might need to gather detailed information relating to the path of a running application 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 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 that is intended for end users, systems administrators and support personnel to view. The text of the message must be clear, concise, and interpretable. Messages are typically localized, meaning that they display 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 for service engineers or developers to use. This trace record might be considerably more complex, verbose, and detailed than a message entry. Localization support is typically not used for trace entries. Trace entries can 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 might be enabled as needed to gather diagnostic information.

WebSphere Application Server provides a message logging and diagnostic trace API that applications can use. This API is based on the stand-alone JRes logging toolkit, which was developed by IBM. The stand-alone JRes 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 JRes logging toolkit provides a limited amount of support, which is typically referred to as *systems management support*, including log file configuration support based on property files.

As designed, the stand-alone JRes logging toolkit does not contain the support that is required for integration into the WebSphere Application Server run time or for use in a Java 2 Platform, Enterprise Edition (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 JRes extensions. The JRes extensions do not modify the interfaces that are introduced by the stand-alone JRes logging toolkit, but provide the appropriate implementation classes. The conceptual structure that is introduced by the stand-alone JRes logging toolkit is described in the following section. It is equally applicable to the JRes extensions.

JRes concepts

The section contains a basic overview of important concepts and constructs that are introduced by the stand-alone JRes logging toolkit. This information is not an exhaustive overview of the capabilities of this logging toolkit, nor is it intended as a detailed discussion of usage or programming paradigms. More detailed information, including code examples, is available in JRes extensions and its subtopics, including in the API documentation for the various interfaces and classes that make up the logging toolkit.

Event types

The stand-alone JRes 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: message loggers and trace loggers. The set of methods on message loggers and trace loggers are different because they provide different functionality. Message loggers create message records only and trace loggers create trace records only. Both types of loggers contain masks that indicate which categories of events the logger processes and which to ignore. Although every JRas logger is defined to contain both a message and trace mask, the message logger uses only the message mask and the trace logger uses the trace mask only. For example, by setting a message logger message mask to the appropriate state, it can be configured to process only error messages and ignore informational and warning messages. Changing the trace mask state of a message logger 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 compares the event type that is specified by the caller to its current mask value. If the specified type passes the mask check, the logger creates an event object to capture the information relating to the event that passed to the logger method. This information can include information, such as the names of the class and method which logs the event, a message, and parameters to log, among others. When the logger creates the event object, it forwards the event to all handlers currently registered with the logger.

Methods that are used within the logging infrastructure do not make calls to the logger method. When an application uses an object that extends a thread class, implements the hashCode method, 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 processes. For example, a message logger might be configured to pass both warning and error events, but a handler attached to the message logger might be configured to pass error events only. 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 can contain multiple formatters, each of which knows 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 are the collection of implementation classes that support JRas integration into the WebSphere Application Server environment.

JRas extensions

The JRas framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

The stand-alone JRas logging toolkit defines interfaces and provides a variety of concrete classes that implement these interfaces. Because the stand-alone JRas logging toolkit is developed as a general purpose toolkit, the implementation classes do not contain the configuration interfaces and methods that are necessary for use in the WebSphere Application Server product. In addition, many of the implementation classes are not written appropriately for use in a Java 2 Platform, Enterprise Edition (J2EE) environment. To overcome these shortcomings, WebSphere Application Server provides the appropriate implementation classes that support 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 that are defined and maintained by the WebSphere Application Server run time. This mode is the default mode of operation and is equivalent to the WebSphere Application Server V4.0 mode of operation.

Stand-alone

In this mode, message and trace records are written solely to stand-alone logs that are 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 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 are defined that support JRas extensions in either a stand-alone or combined mode of operations.

JRas extension classes

WebSphere Application Server provides a base set of implementation classes that are collectively 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.

The JRas framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

The collection of JRas 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 that are 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 that are 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. You cannot directly instantiate message and trace loggers. Obtaining loggers in any manner other than directly from the Manager class is not allowed and directly violates the programming model.

The message and trace logger instances that are obtained from the WebSphere Application Server Manager class are subclasses of the `RASMessageLogger` and `RASTraceLogger` classes that are provided by the stand-alone JRas logging toolkit. The `RASMessageLogger` and `RASTraceLogger` classes define the set of methods that are directly available. Public methods that are introduced by the JRas extensions logger subclasses cannot be called directly by user code because it 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 is ever created within the lifetime of a process. The first call to the Manager class with a particular name results in the logger, which is configured by the Manager class. The Manager class caches a reference to the instance, then returns it to the caller. Subsequent calls to the Manager class that specify the same name result in a returned reference to the cached logger. Separate namespaces are maintained for message and trace loggers. You can use a single name 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 can be used to instrument an entire application. You might determine that having a logger per class is more effective, or the appropriate granularity might be somewhere in between. Partitioning an application into logging domains is determined by the application writer.

The WebSphere Application Server logger classes that are obtained from the Manager class are thread-safe. Although the loggers provided as part of the stand-alone JRas logging toolkit implement the serializable interface, 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.

Personal or individual logger subclasses are not supported 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 is not available to user code. Every logger that is 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. When you remove it, you cannot add the WebSphere Application Server handler again to the logger from which it is 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 in the `com.ibm.websphere.ras` package. All message and trace loggers must be obtained from this Manager class. A reference to the Manager class is obtained by calling the static `Manager.getManager` method. Message loggers are obtained by calling the `createRASMessageLogger` method on the Manager class. 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 supports multiple, unrelated trace loggers to register 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 work.

For example, suppose component A consists of 10 classes. Suppose each class is configured to use a separate trace logger. All 10 trace loggers in the component are registered as members of the same group, for example, `Component_A_Group`. You can 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.

JRas framework (deprecated)

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

Deprecated: The JRas framework described in this topic is deprecated. However, you can achieve similar results using Java logging.

In general, the JRas extensions require you to call the Manager class to obtain a message logger or trace logger. No provision is made for 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. Directly create instances of user-defined handlers and add them to the loggers that are obtained from the Manager class.

The stand-alone JRas logging toolkit provides several handler implementation classes. These handler classes are inappropriate for use in the Java 2 Platform, Enterprise Edition (J2EE) environment. You cannot directly use or subclass any of the Handler classes that are 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 `RASFormatter` interface. All formatters must implement this interface. You can write your own formatter classes that implement the `RASFormatter` interface. You can add these classes to a user-defined handler only. 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 use of these formatter classes is not supported.

Message event types

The stand-alone JRas toolkit defines message event types in the `RASMessageEvent` interface. In addition, the WebSphere Application Server reserves a range of message event types for future use. The `RASMessageEvent` 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 that are retrieved from the Manager class have their message masks set to pass or process all message event types defined in the `RASMessageEvent` interface. 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 is obtained from the Manager class. WebSphere Application Server does not provide any built-in systems management support for managing 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 that are 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. When 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 that are defined in the `RASITraceEvent` interface.

Trace loggers that are retrieved from the Manager class 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, you can change the state of the trace mask for a logger at run-time, using WebSphere Application Server systems management facilities.

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 the handlers that are 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. When 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 processed 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.

JRas programming interfaces for logging (deprecated):

The JRas framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

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. Because 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, locating user-written extensions in a separate package is a good practice. It is your responsibility to restrict access to the user-written extensions appropriately.

Writing a handler

User-written handlers must implement the `RASHandler` interface. The `RASHandler` interface extends the `RASMaskChangeGenerator` interface, which extends the `RASObject` interface. A short discussion of the methods that are 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 API documentation.

RASObject interface

The `RASObject` 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 public `Hashtable getConfig` and public `void setConfig(Hashtable ht)` methods are used to get and set the configuration state. The JRas extensions do not support properties-based configuration. Implement these methods 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 public `String getName` and public `void setName(String name)` methods 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 public `String getDescription` and public `void setDescription(String desc)` methods 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 public `String getGroup` method is provided for use by the `RASManager` interface. Since the JRas extensions provide their own `Manager` class, this method is never called. Implement this as a no-operation.

RASMaskChangeGenerator interface

The `RASMaskChangeGenerator` interface is the interface that defines the implementation methods for filtering of events based on a mask state. 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 because the message logger never generates trace events. Handlers, however, can actively use both mask values. For example, a single handler can handle both message and trace events.

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

- The public `void addMaskChangeListener(RASMaskChangeListener listener)` and public `void removeMaskChangeListener(RASMaskChangeListener listener)` methods are used to add or remove

listeners to the handler. The public Enumeration `getMaskChangeListeners` method returns an enumeration over the list of currently registered listeners. The public void `fireMaskChangedEvent(RASMaskChangeEvent mc)` method 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 mask. The logger 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 mask. For example, the message logger composite mask is formed by *or'ing* the message masks of all handlers that are registered with that logger, then *and'ing* the result with the logger message mask.

All handlers are required to properly implement these methods. In addition, when a user handler is instantiated, the logger that is added must be registered with the handler; use 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. With this process, the logger can dynamically maintain the composite mask.

The `RASMaskChangedEvent` class is defined by the stand-alone JRas logging toolkit. Direct use of that class by user code is supported in this context.

In addition, the `RASIMaskChangeGenerator` interface introduces the concept of caching the names of all message and trace event classes that the implementing object process. The intent of these methods is to support a management program such as a graphical user interface 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.

- The public void `addMessageEventClass(String name)` and public void `removeMessageEventClass(String name)` methods can be called to add or remove a message event class name from the list. The method public Enumeration `getMessageEventClasses` returns an enumeration over the list of message event class names. Similarly, the public void `addTraceEventClass(String name)` and public void `removeTraceEventClass(String name)` methods can be called to add or remove a trace event class name from the list. The public Enumeration `getTraceEventClasses` method returns 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. Because spawning of threads is not supported in the WebSphere Application Server environment, it is expected that handlers do not queue or batch events, although this activity is not expressly prohibited.

- The public int `getMaximumQueueSize()` and public void `setMaximumQueueSize(int size)` methods create `IllegalStateException` exceptions to manage the maximum queue size. The public int `getQueueSize` method is provided to query the actual queue size.
- The public int `getRetryInterval` and public void `setRetryInterval(int interval)` methods support the notion of error retry, which implies some type of queueing.
- The public void `addFormatter(RASIFormatter formatter)`, public void `removeFormatter(RASIFormatter formatter)` and public Enumeration `getFormatters` methods 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 public void `openDevice`, public void `closeDevice` and public void `stop` methods are provided to manage the underlying device that the handler abstracts.
- The public void `logEvent(RASIEvent event)` and public void `writeEvent(RASIEvent event)` methods are provided to pass events to the handler for processing.

Writing a formatter

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

RASIFormatter interface

- The public void setDefault(boolean flag) and public boolean isDefault methods are used by the concrete RASHandler classes that are provided by the stand-alone JRes logging toolkit to determine if a particular formatter is the default formatter. Because 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 public void addEventClass(String name), public void removeEventClass(String name) and public Enumeration getEventClasses methods are provided to determine which event classes a formatter can use to format. You can provide the appropriate implementations.
- The public String format(RASIEvent event) method is called by handler objects and returns a formatted String representation of the event.

Programming model summary

The programming model that is 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 JRes extensions in a manner that does not conform to the following programming guidelines is prohibited.

Note: The JRes framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

You can use the WebSphere Application Server JRes 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 that are provided by the stand-alone JRes logging toolkit directly, unless specifically noted otherwise. Direct usage of those classes is not supported. IBM Support provides no diagnostic aid or bug fixes relating to the direct use of classes that are provided by the stand-alone JRes logging toolkit.
- You must obtain message and trace loggers directly from the Manager class. You cannot directly instantiate loggers.
- You cannot replace the WebSphere Application Server message and trace logger classes.
- You must guarantee that the logger names that are passed to the Manager class are unique, and follow the documented naming constraints. When a logger is obtained from the Manager class, 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 class results in the Manager class creating a logger that is associated with that name. Subsequent calls to the Manager class that specify the same name result in a returned reference to the existing logger.
- The Manager class maintains a hierarchical namespace for loggers. Use a dot-separated, fully qualified class name to identify any logger. Other than dots or periods, logger names cannot contain any punctuation characters, such as an asterisk (*), a comma (,), an equals sign (=), a colon (:), or quotes.
- Group names must comply with the same naming restrictions as logger names.
- The loggers returned from the Manager class are subclasses of the RASMessageLogger and the RASTraceLogger classes that are provided by the stand-alone JRes logging toolkit. You can call any public method that is defined by the RASMessageLogger and RASTraceLogger classes. You cannot call any public method that is 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 cannot use the Handler and Formatter classes that are provided by the stand-alone JRes logging toolkit. User written handlers and formatters must conform to the documented guidelines.

- Loggers that are obtained from the Manager class come with a WebSphere Application Server handler installed. This handler writes message and trace records to logs that are defined by the WebSphere Application Server run time. 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 supported. You are responsible for creating an instance of the handler to add, configuring the handler by setting the handler 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 might get a reference to the handler that is 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 that are obtained in this way. You can remove the WebSphere Application Server handler from the logger by calling the logger removeHandler method, passing in the reference to the WebSphere Application Server handler. When removed, the WebSphere Application Server handler cannot be added again to the logger.
- You can define your own message type. The behavior of user-defined message types and restrictions on their definitions is discussed in JRas framework (deprecated).
- You can define your own message event classes. The use of user-defined message event classes is discussed in JRas framework (deprecated).
- You can define your own trace types. The behavior of user-defined trace types and restrictions on your definitions is discussed in JRas framework (deprecated).
- You can define your own trace event classes. The use of user-defined trace event classes is discussed in JRas framework (deprecated).
- 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 used to manage the predefined types, these updates must not modify the state of any of the bits that correspond 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

The basic JRas message and event types are not the same as those natively recognized by WebSphere Application Server, so the JRas types are mapped onto the types that are native to the runtime environment. You can control the way JRas message and trace events are processed using custom filters and message controls.

Event types

The JRas framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

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

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 that is provided by the stand-alone JRas logging toolkit.

Table 10. Platform message event types. 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

Table 10. Platform message event types (continued). These message types are mapped onto the native message types, as follows.

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

Platform trace event types

The trace event types that are recognized and processed by the WebSphere Application Server run time are defined in the RASITraceEvent interface that is provided by the stand-alone JRas logging toolkit. The RASITraceEvent interface provides a rich and 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, the RASITraceEvent interface provides TYPE_LEVEL1, TYPE_LEVEL2, and TYPE_LEVEL3. The implementations provide support for this set of levels. The levels are hierarchical, enabling level 2 also enables 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, the RASITraceEvent interface 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:

Table 11. WebSphere Application Server native types and JRas RASITraceEvent level types. Mapping WebSphere Application Server trace types to the JRas RASITraceEvent level types.

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

Table 12. WebSphere Application Server native types and JRas RASITraceEvent enumerated types. Mapping WebSphere Application Server trace types to the 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. If you decide to use the non-level types, choose one type from each category and use those types consistently throughout the application, to avoid confusion.

Message and trace parameters

The various message logging and trace method signatures accept the Object, Object[] and Throwable parameter types. WebSphere Application Server processes and formats the various parameter types as follows:

Primitives

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

Object

The `toString` method is called on the object and the resulting `String` is displayed. Implement the `toString` method appropriately for any object that is 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[]` type is provided for the case when more than one parameter is passed to a message logging or trace method. The `toString` method is called on each `Object` in the array. Nested arrays are not handled, that is none of the elements in the `Object` array belong in an array.

Throwable

The stack trace of the `Throwable` type is retrieved and displayed.

Array of primitives

An array of primitive, for example, `byte[]`, `int[]`, is recognized as an `Object`, but is loosely associated by Java code. In general, avoid arrays of primitives, if possible. If arrays of primitives are passed, the results are indeterminate and can 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 needs to 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 problems can result:

- `[B@924586a0b` - This message is deciphered as a byte array at location X. This message is typically returned when an array is passed as a member of an `Object[]` type and results from calling the `toString` method on the `byte[]` type.
- `Illegal trace argument : array of long`. This response 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 problem can occur when a byte array is passed to a method taking a single `Object`. This behavior is subject to change and cannot be relied on.
- `"1" "2"`: The `String` representation of the members of an `int[]` type formed by converting each element to an integer and calling the `toString` method on the integers. This behavior is subject to change and cannot be relied on.
- `[Ljava.lang.Object;@9136fa0b` : An array of objects. Typically this response 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 that is defined in the `RASIMessageEvent` interface.
2. Logging of that message event type must be enabled by the state of the message logger mask.
3. The message event type must pass any filtering criteria that is established by the WebSphere Application Server run-time.

When a WebSphere Application Server logger is obtained from the `Manager` class, the initial setting of the mask forwards 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 mask.

Some editions of the product support user specified message filter levels for a server process. When such a filter level is set, only messages at the specified severity levels are written to WebSphere Application Server. Message types that pass the mask check of the message logger can be filtered out by WebSphere Application Server.

Control tracing

Each edition of the product provides a mechanism for enabling or disabling trace. The various editions can 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 that is defined in the RASITraceEvent interface.
2. Logging of that trace event type must be enabled by the state of the trace logger mask.
3. The trace event type must pass any filtering criteria that is established by the WebSphere Application Server run-time.

When a logger is obtained from the Manager class, 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 is specified. Unlike message loggers, the WebSphere Application Server can dynamically modify the trace mask state of a trace logger. WebSphere Application Server only modifies the portion of the trace logger mask that corresponds to the values that are defined in the RASITraceEvent interface. WebSphere Application Server does not modify undefined bits of the mask that might be in use for user-defined types.

When the dynamic trace enablement feature that is available on some platforms is used, the trace state change is reflected both in the application server run time and the trace mask of the trace logger. If user code programmatically changes the bits in the trace mask corresponding to the values that are defined by in the RASITraceEvent interface, the mask state of the trace logger and the run time state become unsynchronized and unexpected results occur. Therefore, programmatically changing the bits of the mask corresponding to the values that are defined in the RASITraceEvent interface is not supported.

Instrumenting an application with JRas extensions

You can create an application using JRas extensions.

Before you begin

The JRas framework that is described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

About this task

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

Procedure

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

Creating JRas resource bundles and message files

The WebSphere Application Server message logger provides the message and msg methods so the user can log localized messages. In addition, the message logger provides the textMessage method to log messages that are not localized. Applications can use either or both, as appropriate.

Before you begin

The JRas framework that is described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

About this task

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

The `PropertyResourceBundle` class 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 that are displayed on the administrative 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. 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 resource bundles with the application.

Late binding

The application can choose to have the WebSphere Application Server run time localize the message in the process where it displays. 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 administrative console or log viewing program might be run. You must install the `.jar` file in a directory that is part of the extensions class path. In addition, if you forward logs to IBM service, you must also forward the `.jar` file that contains the resource bundles.

To create a resource bundle, perform the following steps.

Procedure

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 number 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 using this approach is not recommended because escaping characters is error prone and confusing. Instead, use a valid separator character that does not display 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 that remain 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. Translate the file into localized versions of the file with language-specific file names for example, the `DefaultMessages.properties` file 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 used to convert the messages into the requested national language and locale.
 4. When a message logger is obtained from the JRes manager, configure the logger to use a particular resource bundle. Messages logged through the message API 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 that is specified by a message key, ensuring that the message is delivered in the correct language.
 5. If the message loggers `msg` method is called, explicitly identify a resource bundle name.

What to do next

The application locates the resource bundle based on the file location relative to any directory in the class path. For instance, if the `DefaultMessages.properties` property resource bundle is 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.

JRes resource bundles:

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

Resource bundle sample

The JRes framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

For this sample, four localizable messages are provided. The properties file is created and the key-value pairs are 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 that are imposed on the values by the Java `MessageFormat` class. For example, apostrophes must be escaped or they cause a problem. Avoid the use of non-portable characters. WebSphere Application Server does not support the use 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 is in the class path. Assume that the properties file is stored in the subdirectory `baseDir` that is not in the class path (`baseDir/subDir1/subDir2/resources`). To allow the messages file to resolve, the `subDir1.subDir2.resources.DefaultMessage` name is used to identify the `PropertyResourceBundle` resource bundle and is passed to the message logger.

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

```
# Contents of the 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 substitution parameters: parm1={0}, parm2 = {1}, parm3={2}
```


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

JRas manager and logger instances

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

Deprecated: The JRas framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

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

In the combined mode, message and trace events are logged to both WebSphere Application Server and user-defined logs.

In the stand-alone mode, message and trace events are logged only to user-defined logs.

Using the message and trace loggers

Regardless of the mode of operation, the use of message and trace loggers is the same.

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

Because trace is normally disabled, guard trace methods 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 you gather parameters to trace.
    if (trcLogger.isLoggable(RASITraceEvent.TYPE_ENTRY_EXIT) {
        // I want to trace three 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. An important state change is
```

```

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

Use JRas operations in integrated mode to send trace events and logging messages to only WebSphere Application Server logs.

Before you begin

The JRas framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

About this task

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

Procedure

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 class and create the loggers. Because loggers are named singletons, you can do this activity in a variety of places. One logical candidate for enterprise beans is the `ejbCreate` method. For example, for the `myTestBean` enterprise bean, 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 that is 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

Use JRas operation in combined mode to output trace data and logging messages to both WebSphere Application Server and user-defined logs.

Before you begin

The JRas framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

About this task

In combined mode, messages and trace are logged to both WebSphere Application Server logs and user-defined logs. The following sample assumes that:

- You wrote a user-defined handler named `SimpleFileHandler` and a user-defined formatter named `SimpleFormatter`.
- You are not using user-defined types or events.

Procedure

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 class, create the loggers, and add the user handlers. Because 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 support this approach. 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 updates 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

You can configure JRas operations to output trace data and logging messages to only user-defined locations.

Before you begin

The JRas framework described in this task and its sub-tasks is deprecated. However, you can achieve similar results using Java logging.

About this task

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.
- You are not using user-defined types of events.

Procedure

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 class, create the loggers, and add the user handlers. Because 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 support this approach. 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 that is 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(RASMessageEvent.DEFAULT_MESSAGE_MASK);
```

Logging Common Base Events in WebSphere Application Server

WebSphere Application Server uses Common Base Events within its basic logging framework. Common Base Events can be created explicitly and then logged through the Java logging API, or can be created implicitly by using the Java logging API directly.

About this task

Attention: Logging Common Base Events is not supported with the High Performance Extensible Logging (HPEL) log and trace mode.

An *event* is a notification from an application or the application server that reports information that is related to a specific problem or situation. Common Base Events provide you with a standard structure for these event notifications, which allow you to correlate events that are received from different applications. Log Common Base Events to capture events from different sources to help you fix a problem within an application environment or to tune system performance.

For Common Base Event creation, the application server environment provides a Common Base Event factory with a content handler that provides both runtime data and template data for Common Base Events.

Procedure

1. Optional: Read about the Common Base Event types and how they are implemented within an application server. Refer to “The Common Base Event in WebSphere Application Server.”
2. Read “Logging Common Base Events in WebSphere Application Server” on page 84.
3. Configure the Common Base Event framework for your application server using one of the following methods:
 - “Logging with Common Base Event API and the Java logging API” on page 72
 - “Generate Common Base Event content with the default event factory” on page 74.

Results

Common Base Events will now be logged according to your configuration. Use these event logs to determine the source of application problems.

The Common Base Event in WebSphere Application Server

The Common Base Event is an XML document that defines a common representation of events that is intended for use by enterprise management and business applications. The Common Base Event defines common fields, the values they can take, and the exact meanings of these values.

An application creates an event object whenever something happens that either needs to be recorded for later analysis or which might require the trigger of additional work. An *event* is a structured notification that reports information that is related to a situation. An event reports three kinds of information:

- The situation: What happened
- The identity of the affected component: For example, the server that shut down
- The identity of the component that is reporting the situation, which might be the same as the affected component

The application that creates the event object is called the *event source*. Event sources can use a common structure for the event. The accepted standard for such a structure is called the *Common Base Event*. The Common Base Event is an XML document that is defined as part of the autonomic computing initiative.

The Common Base Event model is a standard that defines a common representation of events that is intended for use by enterprise management and business applications. This standard, which is developed by the IBM Autonomic Computing Architecture Board, supports encoding of logging, tracing, management, and business events using a common XML-based format. This format makes it possible to correlate different types of events that originate from different applications. For more information about the Common Base Event model, see the Common Base Event specification (*Canonical Situation Data Format: The Common Base Event V1.0.1*). The common event infrastructure currently supports Version 1.0.1 of the specification.

Note:

For WebSphere Application Server Version 8.5, if you delete an application server that was previously deployed with the Common Event Infrastructure (CEI) enabled and you did not uninstall CEI before deleting the server, you must use a different name when creating an application server that you want to deploy with CEI. If you deploy CEI on an application server that was created with the exact same server name as the server that was previously deleted and CEI was not uninstalled, the following error occurs:

```
com.ibm.websphere.management.exception.AdminException: ADMA5026E: No valid target is specified in ObjectName  
WebSphere:cell=targetCell,node=targetNode,server=targetServer for module EventServerMdb.jar+META-INF/ejb-jar.xml
```

If you did not uninstall CEI before deleting the application server, you must ensure that you use a name for the new application server that is different from the name of the server that was previously deployed with the common event infrastructure.

The basic concept behind the Common Base Event model is the *situation*. A situation can be anything that happens anywhere in the computing infrastructure, such as a server shutdown, a disk-drive failure, or a failed user login. The Common Base Event model defines a set of standard situation types that accommodate most of the situations that might arise (for example, StartSituation and CreateSituation).

The Common Base Event contains all of the information that is needed by the consumers to understand the event. This information includes data about the runtime environment, the business environment, and the instance of the application object that created the event.

For complete details on the Common Base Event format, see the XML schema that is included in the Common Base Event specification document, at <http://www.ibm.com/developerworks/autonomic/books/fpy0mst.htm#HDRCBEDDESC> .

Types of problem determination events

Problem determination involves multiple types of data, including at least two different classes of event data, log events, and diagnostic events.

Log events, which are also referred to as *message events*, are typically emitted by components of a business application during normal deployment and operations. Log events might identify problems, but these events are also normally available and emitted while an application and its components are in production mode. The target audience for log and message events is users and administrators of the application and the components that make up the application. Log events are normally the only events available when a problem is first detected, and are typically used during both problem recovery and problem resolution.

Diagnostic events, which are commonly referred to as *trace events*, are used to capture internal diagnostic information about a component, and are usually not emitted or available during normal deployment and operation. The target audience for diagnostic events is the developers of the components that make up the business application. Diagnostic events are typically used when trying to resolve problems within a component, such as a software failure, but are sometimes used to diagnose other problems, especially when the information provided by the log events is not sufficient to resolve the problem. Diagnostic events are typically used when trying to resolve a problem.

A *Common Base Event* is a common structure for an event. It defines common fields, the values that these fields can take, and the exact meanings of these values for an event. Common Base Events are primarily used to represent log events.

Common Base Event structure

A *Common Base Event* is a common structure for an event. It defines common fields, the values that these fields can take, and the exact meanings of these values for an event.

The Common Base Event contains several structural elements. These elements include:

- Common header information
- Component identification, both source and reporter
- Situation information
- Message data
- Extended data
- Context data
- Associated events and association engine

Each of these structural elements has its own embedded elements and attributes.

The following table presents a summary of all the fields in the Common Base Event and their usage requirements for problem determination events.

Table 13. Field name, log events, and base specification. This table shows whether a particular element or attribute is required, recommended, optional, prohibited, or discouraged for log events, and the base specification.

Field name	Log events	Base specification
Version	Required	Required
creationTime	Required	Required
severity	Required	Optional
Msg	Required	Optional
sourceComponentId*	Required	Required
sourceComponentId.location	Required	Required
sourceComponentId.locationType	Required	Required
sourceComponentId.component	Required	Required
sourceComponentId.subComponent	Required	Required
sourceComponentId.componentIdType	Required	Required
sourceComponentId.componentType	Required	Required
sourceComponentId.application	Recommended	Optional
sourceComponentId.instanceId	Recommended	Optional
sourceComponentId.processId	Recommended	Optional
sourceComponentId.threadId	Recommended	Optional
sourceComponentId.executionEnvironment	Optional	Optional
situation*	Required	Required
situation.categoryName	Required	Required
situation.situationType*	Required	Required
situation.situationType.reasoningScope	Required	Required
situation.situationType.(specific Situation Type elements)	Required	Required
msgDataElement*	Recommended	Optional
msgDataElement .msgId	Recommended	Optional
msgDataElement .msgIdType	Recommended	Optional
msgDataElement .msgCatalogId	Recommended	Optional
msgDataElement .msgCatalogTokens	Recommended	Optional
msgDataElement .msgCatalog	Recommended	Optional
msgDataElement .msgCatalogType	Recommended	Optional

Table 13. Field name, log events, and base specification (continued). This table shows whether a particular element or attribute is required, recommended, optional, prohibited, or discouraged for log events, and the base specification.

Field name	Log events	Base specification
msgDataElement .msgLocale	Recommended	Optional
extensionName	Recommended	Optional
localInstanceid	Optional	Optional
globalInstanceid	Optional	Optional
priority	Discouraged	Optional
repeatCount	Optional	Optional
elapsedTime	Optional	Optional
sequenceNumber	Optional	Optional
reporterComponentId*	Optional	Optional
reporterComponentId.location	Required (2)	Required (2)
reporterComponentId.locationType	Required (2)	Required (2)
reporterComponentId.component	Required (2)	Required (2)
reporterComponentId.subComponent	Required (2)	Required (2)
reporterComponentId.componentIdType	Required (2)	Required (2)
reporterComponentId.componentType	Required (2)	Required (2)
reporterComponentId.instanceid	Optional	Optional
reporterComponentId.processId	Optional	Optional
reporterComponentId.threadId	Optional	Optional
reporterComponentId.application	Optional	Optional
reporterComponentId.executionEnvironment	Optional	Optional
extendedDataElements*	Note 3	Optional
contextDataElements*	Note 4	Optional
associatedEvents*	Note 5	Optional

Notes:

- Items followed by an asterisk (*) are elements that consist of sub elements and attributes. The fields in those elements are listed in the table directly following the parent element name.
- Some of the elements are optional, but when included, they include sub elements and attributes that are required. For example, the reporterComponentId element has a ComponentIdentification type. The component attribute in ComponentIdentification is required. Therefore, the reporterComponentId.component attribute is required, but only when the reporterComponentId parent element is included.
- The extendedDataElements element can be included multiple times to supply extended data information. See the Extended data section for more information on required and recommended extended data element values.
- The contextDataElements element can be included multiple times to supply context data information.
- The associatedEvents element can be included multiple times to supply correlation data. No recommended uses of this element exist for the producers of problem determination data, and the use of this element is discouraged.

Common header information:

This topic provides additional information about how to format and use these fields for problem determination events, which can be used to clarify and extend the information provided in the other documents.

The Common Base Event specification [CBE101] provides information on the required format of these fields and the Common Base Event Developer's Guide [CBEBASE] provides general usage guidelines.

The common header information in the Common Base Event includes the following information about an event:

- Version: The version of this Common Base Event
- creationTime: The date and time when the event generated
- Severity and priority: The severity of the condition (situation) that is identified by the event
- extensionName: The type of event that was captured
- localInstancelid and globalInstancelid: Identifiers that can be used to quickly identify a specific event within a set of events
- repeatCount and elapsedTime: Information that supports a system to efficiently report multiple events of the same type, by consolidating those events into a single event
- sequenceNumber: Sequence information that supports a system to order a set of events in other ways than time of capture

severity

All problem determination events must provide an indication as to the relative severity of the condition (situation) being reported by providing appropriate values for the severity field in the Common Base Event. The severity field is required for problem determination events. This field is more restrictive than the base specification for the Common Base Event, which lists this field as optional because effective and efficient problem determination requires the ability to quickly identify the information that is needed to resolve a problem as well as prioritize the problems that need addressing.

Table 14. Severity values. The following values are used for problem determination events:

Value	Severity	Description
10	Information	Log information events, normal conditions, and events that are supplied to clarify operations, for example, state transitions, operational changes. These events typically do not require administrator action or intervention.
20	Harmless	Similar to information events, but are used to capture audit items, such as state transitions or operational changes. These events typically do not require administrator action or intervention.
30	Warning	Warnings typically represent recoverable errors, for example a failure that the system can correct. These events can require administrator action or intervention.

Table 14. Severity values (continued). The following values are used for problem determination events:

Value	Severity	Description
40	Minor	Minor errors describe events that represent an unrecoverable error within a component. The failure affects the component ability to service some requests. The business application can continue to perform its normal functions, but its overall operation might be degraded. These events require administrator action or intervention to address the condition.
50	Critical	Critical errors describe events that represent an unrecoverable error within a component. The failure significantly affects the component ability to service most requests. The business application can continue most, but not all of its normal functions and its overall operation might be degraded. These events require administrator action or intervention to address the condition.
60	Fatal	Fatal errors describe events that represent an unrecoverable error within a component. The failure usually results in the complete failure of the component. The business application can continue some normal functions, but its overall operation might be degraded. These events require administrator action or intervention to address the condition.

msg

Refer to “Message data” on page 68 for information on this attribute.

priority

The use of the priority field is discouraged for problem determination events. The severity field is typically used to communicate and evaluate the importance of problem determination events. Use the priority field to enhance the information that is provided in the severity field, that is, prioritize events of the same severity.

extensionName

The extensionName field is used to communicate the type of event that is reported, for example, what general class of events is being reported. In many cases this field provides an indication of what additional data you can expect with the event, for example, optional data values.

repeatCount

The repeatCount field is valid for problem determination events, but is not typically used or supplied by the event producers. This field is used for data reduction and consolidation by event management and analysis systems.

elapsedTime

The elapsedTime field is valid for problem determination events, but is not typically used or supplied by the event producers. This field is used for data reduction and consolidation by event management and analysis systems.

sequenceNumber

The sequenceNumber field is valid for problem determination events. It is typically used only by event producers when the granularity of the event time stamp (the creationTime field) is not sufficient in ordering events. The sequenceNumber field is typically used to sequence events that have the same time stamp value.

Event management and analysis systems can use the sequenceNumber field for a number of reasons, including providing alternative sequencing, not necessarily based on a time stamp. The recommendations here are provided primarily for event producers.

Component identification for source and reporter:

The component identification fields in the Common Base Event are used to indicate which component in the system is experiencing the condition that is described by the event (the sourceComponentID) and which component emitted the event (the reporterComponentID).

Typically, these components are the same, in which case only the sourceComponentID is supplied. Some notes and scenarios on when to use these two elements in the Common Base Event:

- The sourceComponentID is always used to identify the component experiencing the condition that is described by the event.
- The reporterComponentID is used to identify the component that actually produced and emitted the event. This element is typically used only within events that are emitted by a component that is monitoring another component and providing operational information regarding that component. The monitoring component (for example, a Tivoli® agent or hardware device driver) is identified by the reporterComponentID and the component being monitored (for example, a monitored server or hardware device) is identified by the sourceComponentID.

A potential misuse of the reporterComponentID is to identify a component that provides event conversion or management services for a component, for example, identifying an adapter that transforms the events that are captured by a component into Common Base Event format. The event conversion function is considered an extension of the component and not identified separately.

The information that is used to identify a component in the system is the same, regardless of whether it is the source component or reporter component.

Table 15. Component identification for source and reporter. The information that is used to identify a component in the system is the same, regardless of whether it is the source component or reporter component.

Source component	Reporter component	Description
location locationType	Component location	Identifies the location of the component.
component componentIdType	Component name	Identifies the asset name of the component, as well as the type of component.
subcomponent	Subcomponent name	Identifies a specific part or subcomponent of a component, for example a software module or hardware part.
application	Business application name	Identifies the business application or process the component is a part of and provides services for.
instanceId	Operational instance	Identifies the operational instance of a component, that is the actual running instance of the component.

Table 15. Component identification for source and reporter (continued). The information that is used to identify a component in the system is the same, regardless of whether it is the source component or reporter component.

Source component	Reporter component	Description
processId threadId	Operational instance	Identifies the operational instance of a component within the context of a software operating system, that is the operating system process and thread running when the event was produced.
executionEnvironment	Operational instance Component location	Provides additional information about the operational instance of a component or its location by identifying the name of the environment hosting the operational instance of the component, for example the operating system name for a software application, the application server name for a Java 2 Platform, Enterprise Edition (J2EE) application, or the hardware server type for a hardware part.

The Common Base Event specification [CBE101] provides information on the required format of these fields and the Common Base Event Developer's Guide [CBEBASE] provides general usage guidelines. This section provides additional information about how to format and use some of these fields for problem determination events, which can be used to clarify and extend the information that is provided in the other documents.

Component

The Component field in a problem determination event is used to identify the manageable asset that is associated with the event. A manageable asset is open for interpretation, but a good working definition is a manageable asset represents a hardware or software component that can be separately obtained or developed, deployed, managed, and serviced. Examples of typical component names are:

- IBM eServer™ xSeries® model x330
- IBM WebSphere Application Server version 5.1 (5.1 is the version number)
- The name of an internally developed software application for a component

subComponent

The Subcomponent field in a problem determination event identifies the specific part of a component that is associated with the event. The subcomponent name is typically not a manageable asset, but provides internal diagnostic information when diagnosing an internal defect within a component, that is What part failed? Examples of typical subcomponents and their names are:

- Intel Pentium processor within a server system (Intel Pentium IV Processor)
- the enterprise bean container within a web application server (enterprise bean container)
- the task manager within an operating system (Linux Kernel Task Manager)
- the name of a Java class and method (myclass.mycompany.com or myclass.mycompany.com.methodname).

The format of a subcomponent name is determined by the component, but use the convention shown previously for naming a Java class or the combination of a Java class and method is followed. The subcomponent field is required in the Common Base Event.

componentIdType

The componentIdType field is required by the Common Base Event specification, but provides minimal value for problem determination events. For most problem determination events, it is encouraged to

use the value provided in the application field instead of the componentIdType. The componentIdType field identifies the type of component; the application is identified by the application field.

application

The application field is listed as an optional value within the Common Base Event specification, but provide it within problem determination events whenever it this value is available. The only reason this field is not required for problem determination events is that instances exist where the issuing component might not be aware of the overall business application.

instanceId

The instanceId field is listed as an optional value within the Common Base Event specification, but provide this value within problem determination events whenever it is available.

Always provide the instanceId when a software component is identified and identify the operational instance of the component (for example, which operation instance of an installed software image is actually associated with the event). Provide this value for hardware components when these components support the concept of operational instances.

The format of the supplied value is defined by the component, but must be a value that an analysis system can use (either human or programmatic) to identify the specific running instance of the identified component. Examples include:

- **cell, node, server** name for the IBM WebSphere Application Server
- **deployed EAR file name** for a Java enterprise bean
- **serial number** for a hardware processor

processId

The processId field is listed as an optional value within the Common Base Event specification, but provide this value for problem determination events whenever it is available and applicable. Always provide this value for software-generated events, and identify the operating system process that is associated with the component that is identified in the event. Match the format of the thread ID with the format of the operating system (or other running environment, such as a Java virtual machine). This field is typically not applicable or used for events that are emitted by hardware (for example, firmware).

threadId

The threadId field is listed as an optional value within the Common Base Event specification, but provide this value for problem determination events whenever it is available and applicable. Always provide for software-generated events, and identify the active operating system thread when the event was detected or issued. A notable exception to this recommendation is some operating systems or running environments do not support threads. Match the format of the thread ID with the format of the operating system (or other running environment, such as a Java virtual machine). This field is typically not applicable or used for events that are emitted by hardware (for example, firmware).

executionEnvironment

The executionEnvironment field, when used, identifies the immediate running environment that is used by the component being identified. Some examples are:

- the operating system name when the component is a native software application.
- the operating system/Java virtual machine name when the component is a Java 2 Platform, Standard Edition (J2SE) application.
- the web server name when the component is a servlet.
- the portal server name when the component is a portlet.
- the application server name when the component is an enterprise bean.

The Common Base Event specification [CBE101] provides information on the required format of these fields and the Common Base Event Developer's Guide [CBEBASE] provides general usage guidelines.

Situation information:

The situation information is used to classify the condition that is reported by an event into a common set of situations.

The Common Base Event specification [CBE101] provides information on the set of situations defined for the Common Base Event, with the values and formats that are used to describe these situations. The Common Base Event Developer's Guide [CBEBASE] provides general usage guidelines.

Consider the following points regarding situation information for problem determination events:

- Whenever possible, use the situation categorizations and qualifiers that are described in the base Common Base Event specification. Avoid using your own situation definitions as much as possible.
- Not all messages and logs can be classified using the situation definitions that are supplied in the base Common Base Event specification. You can use the OtherSituation categorization to provide your own situation information, but the recommended course of action for problem determination events is to use the ReportSituation categorization, with reportCategory=Log.
- Warning events can be confusing. A warning event (that is an event with severity=warning) typically indicates a recoverable failure, but the situation settings can be interpreted as unrecoverable failures (for example ConnectSituation, successDisposition=UNSUCCESSFUL). Use the appropriate situation categorization so the severity setting indicates the severity of the situation, that is whether the component recovered from the failure.
- The recommended setting for the reasoningScope value is EXTERNAL for all message events.

Message data:

All problem determination Common Base Events must provide human readable text that describes the specific reported event within the msg field of the Common Base Event.

The text that is associated with events representing actual messages or log entries is expected to be translated and localized. Include the msgDataElement element in the Common Base Event whenever internationalized text is provided in the event. This element provides information about how the message text is created and how to interpret it. This information is particularly invaluable when trying to interpret the event programmatically or when trying to interpret the message independent of the locale or language that is used to format the message text.

Prerequisite: Understand the concepts that are associated with creating internationalized messages. A good source of education on these concepts is provided by the documentation that is associated with internationalization of Java information and the usage of resource bundles within the Java language.

The msgDataElement element in the Common Base Event includes the following information about the value of the msg field that is provided with an event:

- The locale of the supplied message text, which identifies how the locale-independent fields within the message are formatted, as well as the language of the message (msgLocale).
- A locale-independent identifier that is associated with the message that can be used to interpret the message independent of the message language, message locale, and message format (msgId and msgIdType).
- Information on how a translated message is created, including:
 - The identifier that is used to retrieve the message template (msgCatalogId).
 - The name and type of message catalog that are used to retrieve the message template (msgCatalog and msgCatalogType).
 - Any locale-independent information that is inserted into the message template to create the final message (msgCatalogTokens).

The Common Base Event specification [CBE101] provides information on the required format of these fields and the Common Base Event Developer's Guide [CBEBASE] provides general usage guidelines. This section provides additional information about how to format and use these fields for problem determination events.

msg

All message, log, and trace events must provide a human-readable message in the msg field of the Common Base Event. The msg field is required for problem determination events, both log events and diagnostic events. This field is more restrictive than the base specification for the Common Base Event, which lists this field as optional; effective and efficient problem determination requires the ability to quickly identify the reported condition. The format and usage of this message is component-specific, but use the following general guidelines:

- Expect the message text that is supplied with messages and log events to be internationalized.
- Provide the locale of the supplied message text with the msgLocale field in the msgDataElement element of the Common Base Event.
- Provide additional information regarding the format and construction of internationalized messages whenever possible, using the msgDataElement element of the Common Base Event.

msgLocale

Provide the message locale whenever message text is provided within the Common Base Event, as is the case with all problem determination events. The msgLocale field is listed as an optional value within the Common Base Event specification, but provide this information within problem determination events whenever possible. The reason this field is not required for problem determination events is that instances exist where the locale information is not provided or available when formatting the Common Base Event.

msgId and msgIdType

Several companies include a locale-independent identifier within internationalized message text that you can use to interpret the described condition by the message text, independent of the message. For example, most messages issued by IBM software look like IEE890I WTO Buffers in console backup storage = 1024, where a unique, locale-independent identifier IEE890I precedes the translated message text. This identifier provides a way to uniquely detect and identify a message independent of location and language. This detection is invaluable for locale-independent and programmatic analysis.

The msgId field is listed as an optional value within the Common Base Event specification, but it must be provided within problem determination events whenever this identifier is included in the message text. Likewise, the msgIdType field is listed as an optional value within the Common Base Event specification, but it must be provided within problem determination events whenever a value is supplied for msgId. Do not supply these fields when the message text is not translated or localized, for example, for trace events.

msgCatalogId

The msgCatalogId field is listed as an optional value within the Common Base Event specification, but provide this value whenever the Common Base Event includes localized or translated message text, for example when providing problem determination events that represent issued messages or log events. This field is not required for problem determination events because not all problem determination events include translated message text. Some cases exist where the value is not provided or available when formatting the Common Base Event. Do not supply this field when the message text is not translated or localized, for example, for trace events.

msgCatalogTokens

The msgCatalogTokens field is listed as an optional value within the Common Base Event specification, but provide this value whenever the Common Base Event includes localized or translated message text, for example when providing problem determination events that represent issued messages or log events. This field is not required for problem determination events because not all problem determination events include translated message text, and cases exist where the value is not

provided or available when formatting the Common Base Event. This value contains the list of locale-independent values or message tokens that are inserted into the localized message text when creating a translated message.

These values are difficult to extract from a translated message without knowing the translated message template that is used to create the message. Do not supply this field when the message text is not translated or localized

The Common Base Event provides several mechanisms for providing additional data about an event, including this field, extended data elements, and extensions to the schema. Always use the `msgCatalogTokens` field to supply the list of message tokens that is included in the message text associated with an event. These values can also be supplied in other parts of the Common Base Event, but they must be included in this field.

msgCatalog and msgCatalogType

The `msgCatalog` and `msgCatalogType` fields are listed as optional values within the Common Base Event specification, but provide this value whenever the Common Base Event includes localized or translated message text, for example when providing problem determination events that represent issued messages or log events. These fields are not required for problem determination events because not all problem determination events include translated message text, and cases exist where the values are not provided or available when formatting the Common Base Event. Do not complete these fields when the message text has is not translated or localized, for example, for trace events.

Extended data:

The Common Base Event provides several methods for including this additional data, including extending the Common Base Event schema or supplying one or more `ExtendedDataElement` elements within the Common Base Event, which is the preferred approach.

The base information that is included in a Common Base Event might not be sufficient to represent all of the information captured by a component when creating a problem determination event.

Use an `ExtendedDataElement` element to represent a single data item. A Common Base Event can contain more than one of these elements, essentially one for each additional data item. A hint to the number and type of `ExtendedDataElement` elements is supplied by the `extensionName` value, but this information is only a hint. The usage of the attributes in the `ExtendedDataElement` element for problem determination events is the same as those for any other Common Base Event.

Sample Common Base Event instance

This XML document is an example of a Common Base Event instance that is generated by a WebSphere Application Server application.

Use the following example for reference:

```
<CommonBaseEvent creationTime="2004-09-18T04:03:28.484Z"
  globalInstanceId="myhost:1095479647062:1899"
  msg="WSVR0024I: Server server1 stopped"
  severity="10"
  version="1.0.1">
  ... several extendedDataElements for internal use only ...
<sourceComponentId component="com.ibm.ws.runtime.component.ServerCollaborator"
  componentIdType="Unknown"
  executionEnvironment="Windows Vista[x86]#5.0"
  instanceId="myhost/myhost/server1"
  location="myhost"
  locationType="Hostname"
  processId="1095479647062"
  subComponent="Unknown"
  threadId="Alarm : 0"
  componentType="http://www.ibm.com/namespaces/autonomic/WebSphereApplicationServer"/>
```



```

<msgDataElement msgLocale="en_US">
  <msgCatalogTokens value="server1" />
  <msgId>WSVR0024I< /msgId>
  <msgCatalogId>WSVR0024I< /msgCatalogId>
  <msgCatalog>com.ibm.ws.runtime.runtime< /msgCatalog>
</msgDataElement>

<situation categoryName="ReportSituation">
  <situationType xsi:type="ReportSituation" reasoningScope="EXTERNAL" reportCategory="LOG"/>
</situation>

</CommonBaseEvent>

```

A number of extendedDataElement elements in the XML are used by WebSphere Application Server, but are not for application use because these elements might change.

The CommonBaseEvent element defines the Common Base Event instance. This element has a set of attributes that are common for all Common Base Events. This set includes the extensionName attribute, which defines the type or class of the Common Base Event instance, the creation time, severity, and priority.

Nested within the CommonBaseEvent element are elements giving more detail about the situation. The first of these elements is the situation element. This classification is standardized.

The CommonBaseEvent element also includes the sourceComponentId and the (optional) reporterComponentId elements. The sourceComponentId element describes where the situation occurred; the reporterComponentId describes where the situation is detected. If the sourceComponentId and the reporterComponentId elements are the same, the reporterComponentId element is omitted.

The attributes of both the sourceComponentId and the reporterComponentId elements are the same. They identify the component type, name, operating system, and network location. The content of these attributes provides vertical correlation of the stack of IT resources that are active when the Common Base Event is created.

Also included in the CommonBaseEvent element are contextDataElements elements that describe the context in which the situation occurred. This context correlates Common Base Event instances that are part of the same work. This correlation is called *horizontal correlation* because an instance of a particular context type correlates events at the same level of abstraction, for example at the business level, the application level, or at the middleware level.

Extended data elements contain additional data that is used to describe a situation. In this example, an extended data element is added by WebSphere Application Server to describe the Java 2 Platform, Enterprise Edition (J2EE) component that generated the Common Base Event instance and some application data.

Sample Common Base Event template

The content handler uses template information to fill in blanks in the Common Base Event when the Common Base Event complete method is called.

Components that use the WebSphere Application Server event factory home can include a Common Base Event template XML file to provide data to populate Common Base Events. Information that is already supplied in the event is not overridden if the same field is supplied in the template.

The following example illustrates a Common Base Event template:

```

<?xml version="1.0" encoding="UTF-8"?>

<TemplateEvent
  version="1.0.1"

```

```

xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:noNamespaceSchemaLocation="templateEvent.xsd">

<CommonBaseEvent
  <sourceComponentId application="My Application" component="com.ibm.componentX"/>
  <extendedDataElements name="Sample ExtendedDataElement name" type="string">
    <values>Sample ExtendedDataElement value</values>
  </extendedDataElements>
</CommonBaseEvent>

</TemplateEvent>

```

Component identification for problem determination

This topic describes types of problem determination events.

A business application is made up of multiple components. A component can be made up of several internal subcomponents. Consistent application of these concepts is critical for effective problem determination of a business application; all of the parts of the application must use the same concepts and assumptions when creating and formatting events. Use the following definitions and examples when creating Common Base Events for problem determination.

Business application

A business application is the business logic and business data that is used to address a set of specific business requirements. A business application consists of several components of multiple types, combined in a unique manner by an enterprise, to provide the functions and resources that are needed to address those requirements. The primary creator and manager of a business application is the enterprise, and each enterprise or company creates unique business applications. Examples of business applications are the Payroll Application for the ACME Corporation and the Inventory Application for Spacely Sprockets.

Components

A business application is created and managed by the enterprise as a set of components. Components are deployable assets, which are developed either by the enterprise or a vendor, and managed by the enterprise. A component might be created by the enterprise, typically for use within a specific business application. For example, the ACME Corporation might create a set of enterprise beans to represent the business logic that is required by their Payroll Application. A component might also be an asset that is produced by a vendor and acquired by an enterprise. Examples of these components are hardware products, such as IBM eServers or Sun Solaris systems, or software products, such as IBM WebSphere Application Server, Oracle Database Servers.

Subcomponents


A specific component, depending on its complexity, might consist of several subcomponents. For example, the IBM WebSphere Application Server consists of many subcomponents, such as the enterprise bean container and the servlet engine. Subcomponent information is typically used only by the creator of the component to service the component, and as such are not separately deployable or manageable resources in the enterprise. The enterprise might deploy a change or update to a subcomponent, but only upon guidance from the component vendor and as part of the vendor's component. For example, a software fix for the enterprise bean container of the IBM WebSphere Application Server is packaged and deployed as a software update to the IBM WebSphere Application Server. Replacement of the processor in an IBM eServer is deployed as a physical part, but only as a part of the original deployed component, the IBM eServer.

Logging with Common Base Event API and the Java logging API

In cases where the events that are generated by the Java logging API are insufficient to describe the event that needs capturing, you can create Common Base Events with the Common Base Event factory APIs.

Before you begin

When you create a Common Base Event, you can add data to the Common Base Event before it is logged. The following diagram illustrates how application code can create and log Common Base Events:



About this task

WebSphere Application Server is configured to use an event factory that automatically populates WebSphere Application Server-specific information into the Common Base Events that it generates. In general, it is good practice to create events using the WebSphere Application Server default Common Base Event factory because this approach ensures consistency of Common Base Event content across events. However, you can create and use other Common Base Event factories.

Common Base Events are initiated and logged in the following sequence:

1. Application code invokes the `createCommonBaseEvent` method on the `EventFactory` class to create a `CommonBaseEvent`.
2. Application code wraps `CommonBaseEvent` event in a `CommonBaseEventLogRecord` record, and adds event-specific data.
3. Application code calls the `CommonBaseEvent` event `complete` method.
4. The `CommonBaseEvent` event invokes the `ContentHandler` `completeEvent` method.
5. The `ContentHandler` handler adds XML template data to the `CommonBaseEvent` event. Not all `ContentHandler` handlers support templates.
6. The `ContentHandler` handler adds runtime data to the `CommonBaseEvent` event.
7. Application code passes the `CommonBaseEventLogRecord` record to the logger using the `Logger.log` method.
8. Logger passes `CommonBaseEventLogRecord` record to Handlers.
9. Handlers format data and write to the output device.

Procedure

- You can use the default Common Base Event factory to generate content. Read “Generate Common Base Event content with the default event factory” on page 74 for more information.

- If you do not wish to use the default event factory, you can create custom content handlers and event factories.
 1. Create a custom factory home. Read “Creating custom Common Base Event factory homes” on page 78.
 2. Create a custom content handler. Read “Creating custom Common Base Event content handlers” on page 77.

Results

After completing all the above steps you will have a Common Base event based on your configuration settings.

Generate Common Base Event content with the default event factory

A default Common Base Event content handler populates Common Base Events with WebSphere Application Server runtime information. This content handler can also use a Common Base Event template to populate Common Base Events.

The default content handler is used when the server creates `CommonBaseEventLogRecords` as would be the case in the following example:

```
// Get a named logger
Logger logger = Logger.getLogger("com.ibm.someLogger");
// Log to the logger -- implicitly the default content handler
// will be associated with the CommonBaseEvent contained in the
// CommonBaseEventLogRecord.
logger.warning("MSG_KEY_001");
```

To specify a Common Base Event template in the previous case, a `Logger.properties` file would need to be provided with an `eventfactory` entry for `com.ibm.someLogger`. If a valid template is found on the classpath, then the `Logger`'s event factory will use the specified template's content in addition to the WebSphere Application Server runtime information when populating Common Base Events. If the template is not found on the classpath, or is invalid, then the `Logger`'s event factory will only use the WebSphere Application Server runtime information when populating Common Base Events.

The default content handler is also associated with the event factory home supplied in the global event factory context. This is convenient for creating Common Base Events that need to be populated with content similar to that generated from the WebSphere Application Server:

```
// Request the event factory from the global event factory home
EventFactory eventFactory =
    EventFactoryContext.getInstance().getEventFactoryHome().getEventFactory(templateName);

// Create a Common Base Event
CommonBaseEvent commonBaseEvent = eventFactory.createCommonBaseEvent();

// Complete the Common Base Event using content from the template (if specified previously)
// and the server runtime information.
eventFactory.getContentHandler().completeEvent(commonBaseEvent);
```

In the previous example, if the template referenced by `templateName` is found on the classpath, and the template is valid, then the event factory home will return an event factory which uses a content handler that combines the template's content with the WebSphere Application Server runtime information when populating Common Base Events. If the template is not found on the classpath, or is invalid, then the event factory home will return an event factory which uses a content handler that uses only the WebSphere Application Server runtime information when populating Common Base Events.

The default content handler populates Common Base Events in the server environment with the following runtime information:

CommonBaseEvent.globalInstanceId

Value: The unique_record_id

Set this value only if the CommonBaseEvent.globalInstanceId value is null before the completeEvent method is called.

CommonBaseEvent.msg

Value: A localized message that is based on the MsgDataElement element.

Set this value only if the CommonBaseEvent.msg message is null before the completeEvent method is called.

CommonBaseEvent.severity

Value: Set based on the value of level set on the CommonBaseEventLogRecord record, if level >= Level.SEVERE, set to 50; if level >= Level.WARNING, set to 30; the default is set to 10.

Set this value only if the CommonBaseEvent.severity value is null before the completeEvent method is called.

CommonBaseEvent.ComponentIdentification.component

Value: Set based on the LoggerName value that is set on the CommonBaseEventLogRecord record.

Set this value only if the CommonBaseEvent.ComponentIdentification.component is null before the completeEvent method is called.

CommonBaseEvent.ComponentIdentification.componentIdType

Value: "Unknown"

Set this value only if the CommonBaseEvent.ComponentIdentification.componentIdType value is null before the completeEvent method is called.

CommonBaseEvent.ComponentIdentification.executionEnvironment

Value: OSname[OSarch]#OSversion

Set this value only if the CommonBaseEvent.ComponentIdentification.executionEnvironment value is null before the completeEvent method is called.

CommonBaseEvent.ComponentIdentification.instanceId

Value: cellName\nnodeName\nserverName

Set this value only if the CommonBaseEvent.ComponentIdentification.instanceId value is null before the completeEvent method is called. Set only in a server environment because this value is ignored in a client application.

CommonBaseEvent.ComponentIdentification.location

Value: The host name

Set this value only if both the CommonBaseEvent.ComponentIdentification.location and the CommonBaseEvent.ComponentIdentification.locationType values are null before the completeEvent method is called.

CommonBaseEvent.ComponentIdentification.locationType

Value: The host name

Set this value only if both the CommonBaseEvent.ComponentIdentification.location and the CommonBaseEvent.ComponentIdentification.locationType values are null before the completeEvent method is called.

CommonBaseEvent.ComponentIdentification.processId

Value: An internally generated representation of the process number.

Set this value only if the CommonBaseEvent.ComponentIdentification.processId value is null before the completeEvent method is called

CommonBaseEvent.ComponentIdentification.subComponent

Value: Set based on values of the sourceClassName and the sourceMethodName names that are set on the sourceClassName.sourceMethodName name of the CommonBaseEventLogRecord record.

Set this value only if the CommonBaseEvent.ComponentIdentification.subComponent values is null before the completeEvent method is called and both the sourceClassName and the sourceMethodName names are set.

CommonBaseEvent.ComponentIdentification.threadId

Value: Set to the value of the Java Virtual Machine (JVM) thread name.

Set this value only if the CommonBaseEvent.ComponentIdentification.threadId values is null before the completeEvent value is called.

CommonBaseEvent.ComponentIdentification.componentType

Value: <http://www.ibm.com/namespaces/autonomic/WebSphereApplicationServer>

Set this value only if the CommonBaseEvent.ComponentIdentification.componentType values is null before the completeEvent method is called.

CommonBaseEvent.MsgDataElement.msgLocale

Value: Set based on the default locale of the JVM.

Set this value only if the CommonBaseEvent.msg value is null before the completeEvent method is called.

CommonBaseEvent.Situation.categoryName

Value: ReportSituation

Set this value only if the CommonBaseEvent.Situation value is null before the completeEvent method is called.

CommonBaseEvent.Situation.situationType.type

Value: ReportSituation

Set this value only if the CommonBaseEvent.Situation value is null before the completeEvent method is called.

CommonBaseEvent.Situation.situationType.reasoningScope

Value: EXTERNAL

Set this value only if the CommonBaseEvent.Situation value is null before the completeEvent method is called.

CommonBaseEvent.Situation.situationType.reportCategory

Value: LOG

Set this value only if the CommonBaseEvent.Situation value is null before the completeEvent method is called.

The sourceComponentIdentification value is populated if no reporterComponentIdentification ID exists when the completeEvent method is invoked on the content handler. Otherwise, the reporterComponentIdentification ID is populated instead.

Common Base Event content handler

Content handlers populate data into Common Base Events when the Common Base Event complete method is invoked. You can associate content handlers with Common Base Event templates, which provide default information to transfer into each Common Base Event.

Content handlers might also provide any other information that is relevant to completing the population of the Common Base Event, such as appropriate runtime defaults. The use of content handlers ensures consistency of field use in the Common Base Event within a component or within a set of components that share the same runtime. For example, some content handlers support the specification of a template. If

used consistently across a component, this template ensures that all events for that component have the same template information filled in. Similarly, some content handlers can also supply runtime information to their associated Common Base Events. If consistently used throughout the entire runtime, runtime information ensures that all events use runtime data in a similar way.

The event factory home that is used in the WebSphere Application Server runtime is associated with a content handler that both reads from a template, and supplies runtime data. Have components use Event Factories that are obtained from this event factory home with their own templates, to produce consistency between application events and server events.

More details can be found in “Creating custom Common Base Event content handlers” or the API documentation for `org.eclipse.hyades.logging.events.cbe.ContentHandler` at <http://www.eclipse.org/tptp/index.html>.

Creating custom Common Base Event content handlers

Create a custom Common Base Event content handler or template to automate configuration or values for specific events.

Before you begin

A *content handler* is an object that automatically sets the property values of each event based on any arbitrary policies that you want to use.

The following content handler classes were added to WebSphere Application Server to facilitate the use of the Common Base Event infrastructure:

Class name	Description
<code>WsContentHandlerImpl</code>	This provides an implementation of <code>org.eclipse.hyades.logging.events.cbe.ContentHandler</code> specifically for use in the WebSphere Application Server environment. This content handler completes Common Base Events using information from the WebSphere Application Server runtime, and it uses the same content handler as is used internally by the WebSphere Application Server when completing Common Base Events for logging.
<code>WsTemplateContentHandlerImpl</code>	This provides the same function as <code>WsContentHandlerImpl</code> , but it extends the <code>org.eclipse.hyades.logging.events.cbe.impl.TemplateContentHandlerImpl</code> class to enable the use of a Common Base Event template. Template content takes precedence in cases where the template data specifies values for the same Common Base Event fields as does the <code>WsContentHandlerImpl</code> .

About this task

In some situations, you might want some event property data set automatically for every event that you create. This automation is a way to fill in certain standard values that do not change, such as the application name, or to set some properties based on information that is available from the runtime environment, like creation time or thread information. You can set property data automatically by creating a content handler.

Procedure

- Use the following code sample to implement the `CustomContentHandler` class:

```
public class CustomContentHandler extends WsContentHandlerImpl {

    public CustomContentHandler() {
        super();
        // TODO Custom initialization code goes here
    }
}
```

```

}

public void completeEvent(CommonBaseEvent cbe) throws CompletionException {
    // following code will add WAS content to the Content Base Event
    super.completeEvent(cbe);
    // TODO Custom content can be added to the Content Base Event here
}
}

```

- The following shows how to implement the CustomTemplateContentHandler class:

```

public class CustomTemplateContentHandler extends WsTemplateContentHandlerImpl {

    public CustomTemplateContentHandler() {
        super();
        // TODO Custom initialization code goes here
    }

    public void completeEvent(CommonBaseEvent cbe) throws CompletionException {
        // following code will add WAS content to the Content Base Event
        super.completeEvent(cbe);
        // TODO Custom content can be added to the Content Base Event here
    }
}

```

Results

You now have a content handler or a custom content handler template based on the settings that you specified.

Common Base Event factory home

Event Factory homes provide Event Factory instantiation that is based on a unique factory name.

Event factory home implementations are tightly coupled with content handlers that are used to populate Common Base Events with template or default data. Event factory instances are maintained by the associated event factory home, based on their unique name. For example, when application code requests a named event factory, the newly created Event Factory instance is returned and persisted for future requests for that named event factory. An abstract event factory home class provides the implementation for the APIs in the event factory home interface. Implementers extend the abstract event factory home class and implement the createContentHandler API to create a typed content handler that is based on the type of event factory home implementation.

In WebSphere Application Server, the default event factory home that is obtained with a call to EventFactoryContext.getInstance.getEventFactoryHome method is associated with a ContentHandler handler capable of supplying both event template information, as well as WebSphere Application Server runtime default information.

More details can be found in the API documentation for `org.eclipse.hyades.logging.events.cbe.EventFactoryHome` at www.eclipse.org/hyades.

Creating custom Common Base Event factory homes

Use custom Common Base Event factory homes to control configuration and implementation of unique event factories.

Before you begin

Event factory homes create and provide homes for Event Factory instances. Each event factory home has a content handler. This content handler is assigned to every event factory the event factory home creates. In turn, when a Common Base Event is created, the content handler from the event factory is assigned to it. Event factory instances are maintained by the associated event factory home, based on their unique

name. For example, when application code requests a named event factory, the newly created event factory instance is returned and persisted for future requests for that named event factory.

The following classes were added to facilitate the use of event eactory homes for logging Common Base Events:

Class name	Description
WsEventFactoryHomeImpl	This class extends the org.eclipse.hyades.logging.events.cbe.impl.AbstractEventFactoryHome class. This event factory home returns event factory instances associated with the WsContentHandlerImpl content handler. The WsContentHandlerImpl is the content handler used by the WebSphere Application Server by default when no event factory template is in use.
WsTemplateEventFactoryHomeImpl	This class extends the org.eclipse.hyades.logging.events.cbe.impl.EventXMLFileEventFactoryHomeImpl class. This event factory home returns event factory instances associated with the WsTemplateContentHandlerImpl Content Handler. The WsTemplateContentHandlerImpl is the content handler used by the WebSphere Application Server when an Event Factory template is required.

About this task

Custom event factory homes support the use of Common Base Event for logging in WebSphere Application Server and make logging easy and consistent between the WebSphere Application Server runtime and the exploiters of this API. The CustomEventFactoryHome and CustomTemplateEventFactoryHome classes will be used to obtain an event factory. These classes are there to make sure the correct content handler is being used with a particular event factory. The CustomEventFactoryHelper class is an example of how the infrastructure provider can hide the factory selection details from infrastructure users, using their own set of parameters to decide which the appropriate event factory is.

Procedure

- The following code samples provide examples of how to implement and use the CustomEventFactoryHome class.

- Implementation of the CustomEventFactoryHome class is as follows:

```
public class CustomEventFactoryHome extends AbstractEventFactoryHome {

    public CustomEventFactoryHome() {
        super();
        // TODO Custom intialization code goes here
    }

    public ContentHandler createContentHandler(String arg0) {
        // Always use custom content handler
        return resolveContentHandler();
    }

    public ContentHandler resolveContentHandler() {
        // Always use custom content handler
        return new CustomContentHandler();
    }
}
```

- The following is an example of how to use the CustomEventFactoryHome class:

```
// get the event factory
EventFactory eventFactory=(new CustomEventFactoryHome()).getEventFactory("XYZ");
// create an event - call appropriate method
eventFactory.createCommonBaseEvent();
// log event ...
```

- For the CustomTemplateEventFactoryHome class you can use the following code for implementation and use:

1. Implement the CustomTemplateEventFactoryHome class by using this code:

```
public class CustomTemplateEventFactoryHome extends
    EventXMLFileEventFactoryHomeImpl {

    public CustomTemplateEventFactoryHome() {
        super();
        // TODO Custom initialization code goes here
    }

    public ContentHandler createContentHandler(String arg0) {
        // Always use custom content handler
        return resolveContentHandler();
    }

    public ContentHandler resolveContentHandler() {
        // Always use custom content handler
        return new CustomTemplateContentHandler();
    }
}
```

2. Use the CustomTemplateEventFactoryHome class by following this sample code:

```
// get the event factory
EventFactory eventFactory=(new
    CustomTemplateEventFactoryHome()).getEventFactory("XYZ");
// create an event - call appropriate method
eventFactory.createCommonBaseEvent();
// log event ...
```

- The CustomEventFactoryHelper class can be implemented and used by following the code below:

1. Implement the custom CustomEventFactoryHelper class using this code:

```
public class CustomTemplateEventFactoryHome extends
    EventXMLFileEventFactoryHomeImpl {

    public CustomTemplateEventFactoryHome() {
        super();
        // TODO Custom initialization code goes here
    }

    public ContentHandler createContentHandler(String arg0) {
        // Always use custom content handler
        return resolveContentHandler();
    }

    public ContentHandler resolveContentHandler() {
        // Always use custom content handler
        return new CustomTemplateContentHandler();
    }
}
```

Figure 4 CustomTemplateEventFactoryHome class

```
public class CustomEventFactoryHelper {
    // name of the event factory to use
    public static final String FACTORY_NAME="XYZ";

    public static EventFactory getEventFactory(String param1, String param2) {
        EventFactory factory=null;
        switch (resolveFactory(param1,param2)) {
            case 1:
                factory=(new CustomEventFactoryHome()).getEventFactory(FACTORY_NAME);
                break;
            case 2:
                factory=(new
                    CustomTemplateEventFactoryHome()).getEventFactory(FACTORY_NAME);
                break;
        }
    }
}
```

```

default:
    // Add default for event factory
    break;
}
return factory;
}

private static int resolveFactory(String param1, String param2) {
    int factory=0;
    // Add code here to resolve which factory to use
    return factory;
}
}

```

2. To use the CustomEventFactoryHelper class, use the following code:

```

// get the event factory
EventFactory eventFactory=
    CustomEventFactoryHelper.getEventFactory("param1","param2","param3");
// create an event - call appropriate method
eventFactory.createCommonBaseEvent();
// log event ...

```

Results

Use the information provided here to implement a custom content factory home and the associated classes based on the settings that you specify.

Common Base Event factory context

The event factory context provides a service to look up event factory homes. Retrieve the event factory context using a call to the EventFactoryContext.getInstance method.

Using this class, you can look up the event factory homes by name, and avoid the need to include the typed home in code. The EventFactoryHome name must be located on the class path to be found. The EventFactoryContext context also stores an EventFactoryHome name as a default, which can be obtained with a call to the EventFactoryContext.getInstance.getEventFactoryHome method.

In WebSphere Application Server, the EventFactoryContext context is configured with a default EventFactoryHome name which is associated to a ContentHandler handler that is capable of supplying both event template information, as well as WebSphere Application Server runtime default information.

More details can be found in the API documentation for `org.eclipse.hyades.logging.events.cbe.EventFactory` at www.eclipse.org/hyades.

Common Base Event factory

Use event factories to create Common Base Events and complete event properties with associated content handlers.

Content handlers populate data into Common Base Events when the Common Base Event invokes the complete method. All event properties set by the application code have priority over all properties that are specified by the content handler. Event factory implementations are tightly coupled with the content handler instance, which is associated with the event factory when the event factory is instantiated. Factory instances can be retrieved only from their associated event factory home. Event factory instances are retrieved and maintained based on unique names. Event factory names are hierarchical; they are represented using the standard Java dot-delimited, name-space naming conventions.

More details can be found in the API documentation for `org.eclipse.hyades.logging.events.cbe.EventFactory` at www.eclipse.org/hyades.

java.util.logging -- Java logging programming interface

The `java.util.logging.Logger` class provides a variety of methods with which data can be logged.

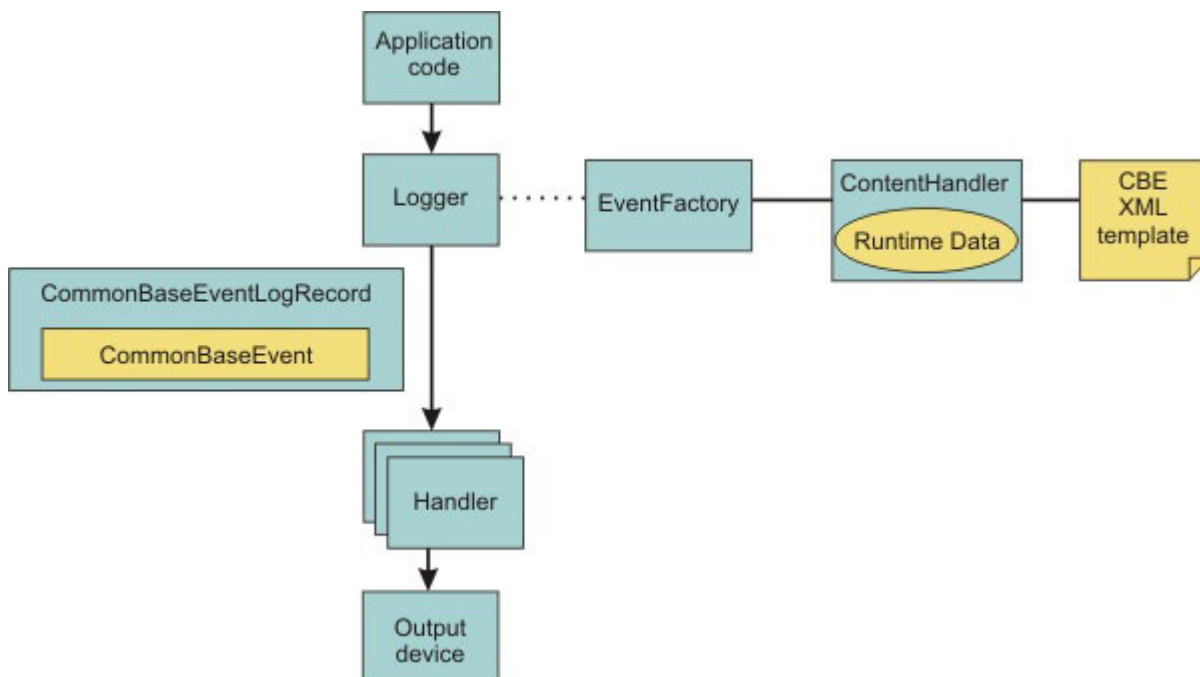
In the WebSphere Application Server, when using basic log and trace mode, the Java logging API (`java.util.logging`) automatically creates Common Base Events for events that are logged at the `WsLevel.DETAIL` level or later (including `WsLevel.DETAIL`, `Level.CONFIG`, `Level.INFO`, `WsLevel.AUDIT`, `Level.WARNING`, `Level.SEVERE`, and `WsLevel.FATAL`). These Common Base Events are created using the event factory that is associated with the logger to which the message is logged. If no event factory is specified, WebSphere Application Server uses a default event factory which automatically fills in WebSphere Application Server-specific information.

The WebSphere Application Server uses a special implementation of the `java.util.logging.Logger` class that automatically creates Common Base Events for the following methods:

- `config`
- `info`
- `warning`
- `severe`
- `log`: All variants except `log(LogRecord)` when used with the `WsLevel.DETAIL` level or more severe levels
- `logp`: When used with the `WsLevel.DETAIL` level or more severe levels
- `logrb`: When used with the `WsLevel.DETAIL` level or more severe levels

The WebSphere Application Server logger implementation is used only for named loggers for example, loggers that are instantiated with calls, such as `Logger.getLogger("com.xyz.SomeLoggerName")`. Loggers instantiated with calls to the `Logger.getAnonymousLogger` and `Logger.getLogger`, or `Logger.global` methods do not use the WebSphere Application Server implementation, and do not automatically create Common Base Events for logging requests made to them. Log records that are logged directly with the `Logger.log(LogRecord)` method are not automatically converted by WebSphere Application Server loggers into Common Base Events.

The following diagram illustrates how application code can log Common Base Events:



The Java logging API processing of named loggers and message-level events proceeds as follows:

1. Application code invokes the named logger (WsLevel.DETAIL or later) with event-specific data.
2. The logger creates a Common Base Event using the createCommonBaseEvent method on the event factory that is associated with the logger.
3. The logger creates a Common Base Event using the event factory associated to the logger.
4. The logger wraps the common base event in a CommonBaseEventLogRecord record, and adds event-specific data.
5. The logger calls the Common Base Event complete method.
6. The Common Base Event invokes the ContentHandler completeEvent method.
7. The content handler adds XML template data to the Common Base Event (including for example, the component name). Not all content handlers support templates.
8. The content handler adds runtime data to the Common Base Event (including for example, the current thread name).
9. The logger passes the CommonBaseEventLogRecord record to the handlers.
10. The handlers format data and write to the output device.

Logger.properties file

Use the `Logger.properties` file to set logger attributes for your component.

The properties file is loaded the first time the `Logger.getLogger(loggername)` method is called within an application. The `Logger.properties` file must be either on the WebSphere Application Server class path, or the context class path.

The logging subsystem uses Common Base Events to represent all the messages in the WebSphere Application Server `activity.log` file. You can specify your own event factory template to be used with your loggers. Use the `eventfactory` property in your `Logger.properties` file. See “Sample Common Base Event template” on page 71 for details on the Common Base Event template.

By convention, the name of the event factory template file should be the fully qualified package name of the package using the template. The name of the file must end with the `.event.xml` extension. For example, a valid event factory template file name for the `com.abc.somepackage` package is:

```
com.abc.somepackage.event.xml
```

When you specify the property value for the `eventfactory` property in the `Logger.properties` file, include the full path name with no leading slash relative to the root of your class path entry. Do not include the `.event.xml` extension.

For example, if the template files from the previous example are located in the `com/abc/templates` directory, the valid value for the `eventfactory` property is:

```
com/abc/templates/com.abc.somepackage
```

Finally, if this event factory template file is used by the `com.abc.somepackage.SomeClass` logger, then the following entry will appear in the `Logger.properties` file:

```
com.abc.somepackage.SomeClass.eventfactory=com/abc/templates/com.abc.somepackage
```

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Logging Common Base Events in WebSphere Application Server

The following practices ensure consistent use of Common Base Events within your components, and between your components and WebSphere Application Server components.

Follow these guidelines:

- Use a different logger for each component. Sharing loggers across components gets in the way of associating loggers with component-specific information.
- Associate loggers with event templates that specify source component identification. This association ensures that the source of all events created with the logger is properly identified.
- Use the same template for directly created Common Base Events (events created using the Common Base Event factories) and indirectly created Common Base Events (events created using the Java logging API) within the same component.
- Avoid calling the complete method on Common Base Events until you are finished adding data to the Common Base Event and are ready to log it. This approach ensures that any decisions made by the content handler based on data already in the event are made using the final data.

The following sample `Logger.properties` file entry demonstrates how to associate the `com.ibm.componentX` logger with the `com.ibm.componentX` event factory:

```
com.ibm.componentX.eventfactory=com.ibm.componentX
```

The following sample code demonstrates the use of the same event factory setting for direct (Part 1) and indirect (Part 2) Common Base Event logging:

```
<?xml version="1.0" encoding="UTF-8"?>

<TemplateEvent>
  version="1.0.1"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:noNamespaceSchemaLocation="templateEvent.xsd">

  <CommonBaseEvent>
    <sourceComponentId application="My application" component="com.ibm.componentX"/>
    <extendedDataElements CommonBaseEventname="Sample ExtendedDataElement name" type="string">
      <values>Sample ExtendedDataElement value</values>
    </extendedDataElements>
  </CommonBaseEvent>

</TemplateEvent>
```

Showlog commands for Common Base Events


The `showlog` command converts the service log from binary format into plain text.

Purpose

These `showlog` commands to produce output in Common Base Event XML format.

-  `showlog -format CBE-XML-1.0.1 filename`

where:

 `filename`
Is the service log file name.

For examples of `showlog` scripts, see [Viewing the service log](#).

Chapter 5. Configuring Java logging using the administrative console

Java logging provides a standard logging API for your applications. Before applications can log diagnostic information, you need to specify how you want the server to handle log output and what level of logging you require.

About this task

Developing, deploying and maintaining applications are complex tasks. When an application encounters an unexpected condition, it might not be able to complete a requested operation. You might want the application to inform the administrator that the operation failed and tell the administrator why the operation failed. This information enables the administrator to take the proper corrective action. Application developers might need to gather detailed information that relates to the path of a running application 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 *logging* and *tracing*. For more information read “Java logging” on page 28.

Using the administrative console, you can:

- Enable or disable a particular log, specify where log files are stored and how many log files are kept.
- Specify the level of detail in a log, and specify a format for log output.
- Set a log level for each logger.

You can change the log configuration statically or dynamically. Static configuration changes affect applications when you start or restart the application server. Dynamic or run time configuration changes apply immediately.

When a logger is created, the level value for that logger is set from the configuration data. If no configuration data is available for a particular logger name, the level for that logger is obtained from the parent of the logger. If no configuration data exists for the parent logger, the parent of that logger is checked, and so on up the tree, until a logger with a non-null level value is found. When you change the level of a logger, the change is propagated to the children of the logger, which recursively propagates the change to their children, as necessary.

Procedure

1. Set the logging levels for your logs:
 - a. In the navigation pane, click **Servers > Server Types > WebSphere application servers**.
 - b. Click the name of the server that you want to work with.
 - c. Under Troubleshooting, click **Logs and Trace**.
 - d. Click **Change Log Detail levels**.
 - e. To make a static change to the configuration, click the **Configuration** tab. A list of well-known components, packages, and groups is displayed. To change the configuration dynamically, click the **Runtime** tab. The list of components, packages, and groups displays all the components that are currently registered on the running server.
 - f. Select a component, package, or group to set a logging level.
 - g. [High Performance Extensible Logging] Select whether or not you want to disable the logging and tracing of potentially sensitive data.
 - h. Click **Apply**.
 - i. Click **OK**.
2. To have static configuration changes take effect, stop then restart the application server.

Log streams and expected output

Investigating the logging and tracing output for the application server is an excellent way to observe performance, diagnose problems, and gain a general understanding of how the application server is working within your environment. The expected output locations for logging and trace information can be different depending on the operating system on which the application server is running.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

IBM i

Table 16. Log and trace output for distributed and IBM i. This table lists the expected output for logging and tracing output streams for distributed and IBM i operating systems

Log or trace stream	Expected output - basic log and trace mode	Expected output - HPEL log and trace mode
Java trace	Trace.log. You must have tracing enabled.	<ul style="list-style-type: none">• HPEL trace repository.• HPEL TextLog*.log, when you enable this log.
Java logs	<ul style="list-style-type: none">• Trace.log, when you enable tracing.• Activity.log. The Activity.log file is also directed to the Showlog output.• SystemOut.log	<ul style="list-style-type: none">• HPEL log repository• HPEL TextLog*.log, when you enable this log.
System.out stream	<ul style="list-style-type: none">• Trace.log• SystemOut.log	<ul style="list-style-type: none">• HPEL log repository• HPEL TextLog*.log, when you enable this log.
System.err stream	SystemErr.log	<ul style="list-style-type: none">• HPEL log repository• HPEL TextLog*.log, when you enable this log
cout (the C or C++ output stream)	Native stdout	Native stdout
cerr (the C or C++ error stream)	Native stderr	Native stderr





Figure 1. Distributed and IBM i - Basic log and trace mode. Distributed and IBM i - Basic log and trace mode

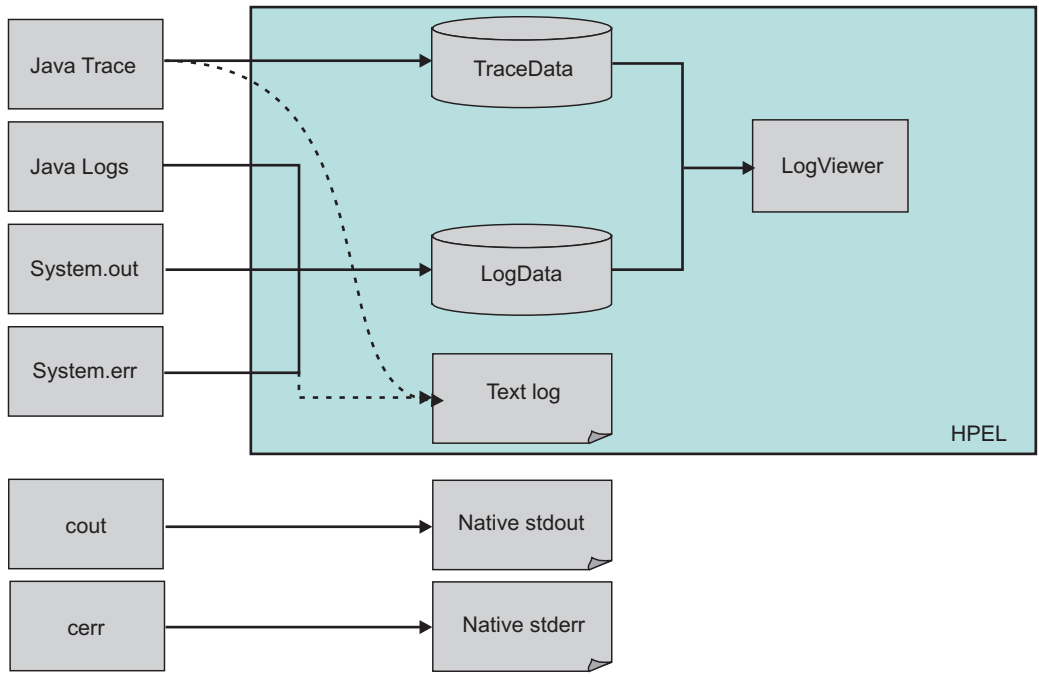


Figure 2. Distributed and IBM i - HPEL log and trace mode. Distributed and IBM i - HPEL log and trace mode

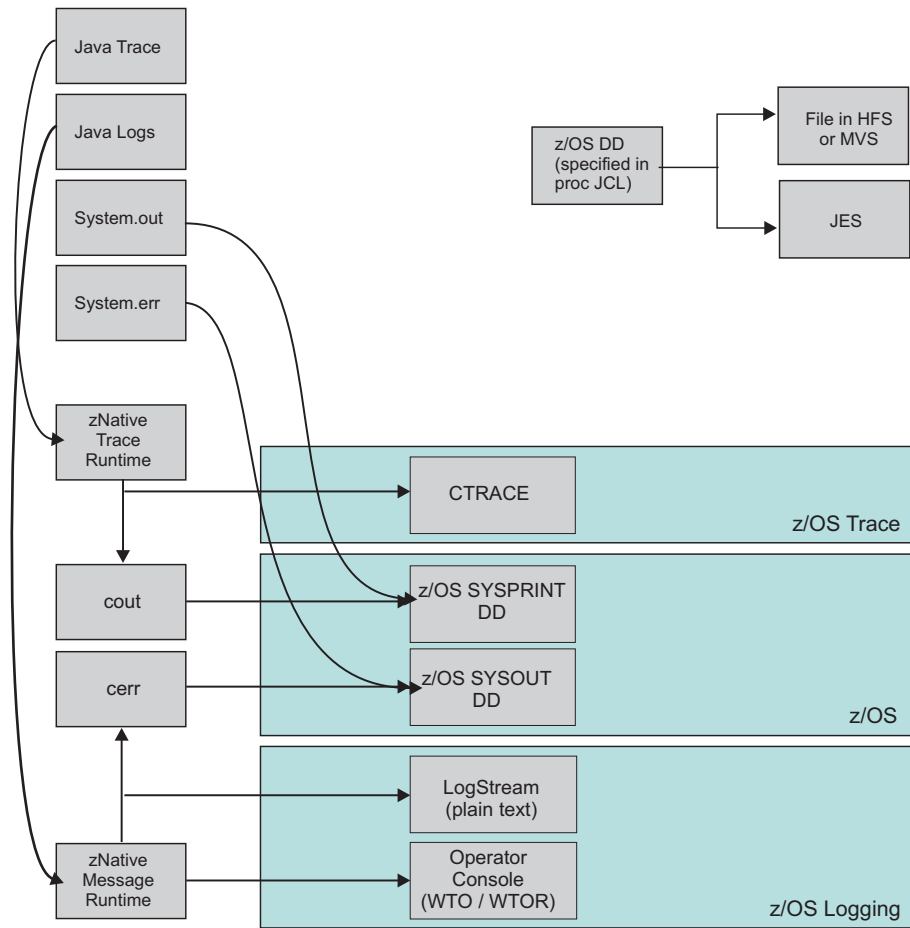


Figure 3. z/OS - Traditional log and trace mode. z/OS - Traditional log and trace mode




Figure 4. z/OS - HPEL log and trace mode. z/OS - HPEL log and trace mode

Log level settings

Use this topic to configure and manage log level settings.

To view this administrative console page, click **Troubleshooting > Logs and trace > *server_name* > Change log detail levels**

Using log levels you can control which events are processed by Java logging. When you change the level for a logger, the change is propagated to the children of the logger.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Disable logging and tracing of potentially sensitive data

The application server has a list of loggers which are known to potentially write sensitive information to the log and trace when enabled. For example, enabling certain HTTP related loggers at FINEST level may result in confidential user-specified information from HTTP requests being stored in the trace files. If you want the server to avoid enabling these loggers at levels which are known to be used for potentially sensitive information, check the **Disable logging and**

tracing of potentially sensitive data check box. When the server is started, or when the log detail level specification is modified at run time, the server will compare the list of loggers and levels specified in the log detail level specification to the list of loggers and levels in the sensitive logger list, and will update the log detail level specification as needed.

Change Log Detail Levels

Enter a log detail level that specifies the components, packages, or groups to trace. The log detail level string must conform to the specific grammar described in this topic. You can enter the log detail level string directly, or generate it using the graphical trace interface.

If you select the **Configuration** tab, and expand **Components and Groups**, a static list of well-known components, packages, and groups is displayed. This list might not be exhaustive.

If you select the **Runtime** tab, and expand **Components and Groups**, the list of components, packages, and groups are displayed with all the components that are registered on the running application server and in the static list.

The format of the log detail level specification is:

```
<component> = <level>
```

where *<component>* is the component for which to set a log detail level, and *<level>* is one of the valid logger levels (off, fatal, severe, warning, audit, info, config, detail, fine, finer, finest, all). Separate multiple log detail level specifications with colons (:).

Components correspond to Java packages and classes, or to collections of Java packages. Use an asterisk (*) as a wildcard to indicate components that include all the classes in all the packages that are contained by the specified component. For example:

- * Specifies all traceable code running in the application server, including the product system code and customer code.

com.ibm.ws.*

Specifies all classes with the package name beginning with com.ibm.ws.

com.ibm.ws.classloader.JarClassLoader

Specifies the JarClassLoader class only.

An error can occur when setting a log detail level specification from the administrative console if selections are made from both the Groups and Components lists. In some cases, the selection made from one list is lost when adding a selection from the other list. To work around this problem, enter the log detail level specification directly into the log detail level entry field.

Select a component or group to set a log detail level. The table following lists the valid levels for application servers at WebSphere Application Server Version 6 and later.

Note: Logging level values are case-sensitive and begin with a lowercase letter.

Table 17. Valid logging levels. The following table lists the valid levels for application servers at WebSphere Application Server Version 6 and later.

Version 6 and later logging level	Content / Significance
off	Logging is turned off.
fatal	Task cannot continue and component, application, and server cannot function.
severe	Task cannot continue but component, application, and server can still function. This level can also indicate an impending unrecoverable error.
warning	Potential error or impending error. This level can also indicate a progressive failure (for example, the potential leaking of resources).

Table 17. Valid logging levels (continued). The following table lists the valid levels for application servers at WebSphere Application Server Version 6 and later.

Version 6 and later logging level	Content / Significance
audit	Significant event affecting server state or resources
info	General information outlining overall task progress
config	Configuration change or status
detail	General information detailing subtask progress
fine	Trace information - General trace + method entry, exit, and return values
finer	Trace information - Detailed trace
finest	Trace information - A more detailed trace that includes all the detail that is needed to debug problems
all	All events are logged. If you create custom levels, All includes those levels, and can provide a more detailed trace than finest.

When you enable a logging level in Version 6.0 and later, you are also enabling all of the levels with higher severity. For example, if you set the logging level to warning on your Version 6.x application server, then warning, severe and fatal events are processed.

[Basic mode logging] Trace information, which are events at the Fine, Finer and Finest levels, can be written only to the trace log. Therefore, if you do not enable diagnostic trace, setting the log detail level to Fine, Finer, or Finest will not have an effect on the data that is logged.

Correlation

Specify correlation settings that you would like to enable. Select the **Enable log and trace correlation** check box to enable correlation for the application server. Clear the **Enable log and trace correlation** check box to disable correlation for the application server. Select **Include request IDs in log and trace records**, **Include request IDs in log and trace records and correlation log records**, or **Include request IDs in log and trace records, create correlation log records, and capture data snapshots**, as appropriate.

Note: Enable XCT to include request IDs in log and trace files when you want to see which log and trace entries, in all threads and application server processes, are related to the same request. Request IDs are only recorded when using HPEL log and trace mode and can be seen or used for filtering using the logViewer command.

Note: Enable XCT to create correlation log records when you want to log how requests branch between threads and processes, and see extra information about each request. Enabling XCT to create correlation log records might have a significant performance impact on your system, so is best suited to test and development environments.

Note: Enable XCT to capture data snapshots when you want to store entire request and response bodies to the file system. Enabling XCT to capture data snapshots might have a significant performance impact on your system, so is best suited to test and development environments. XCT captures data snapshots for message requests and responses handled by the SIBus.

Note: Data snapshots are captured and written to the `$SERVER_LOG_ROOT/snapdata` directory. The application server does not automatically clean up files from this directory. You will need to delete the files from this directory periodically when data snapshot capturing is enabled. Data snapshots store entire request and response contents and may include sensitive information. This option might not be appropriate for use in production environments.

Save runtime changes to configuration as well

Specifies that changes are made to both the dynamic state of the running server, and the server configuration, which takes effect on the next restart. If this check box is not selected, the server does not copy the settings into the server configuration.

Changing the message IDs used in log files

You can change the default format for message IDs in server logs by setting the `com.ibm.websphere.logging.messageId.version` system property.

Before you begin

Note: Beginning with WebSphere Application Server Version 6.0, logging files are formatted according to a standardized system. However, the default runtime behavior is still configured to use the older format. In new releases of WebSphere Application Server, the message IDs that are written to log files will be changed to ensure they do not conflict with other IBM products. The default runtime behavior is still configured to use the older message IDs, deprecated in Version 8.5.

As a result of the default runtime behavior, you might see a mixture of messages that use 4-letter message prefixes and 5-letter message prefixes. The information in this topic explains how to change your configuration so that the messages consistently show with 5-letter message prefixes. The default behavior has not changed to minimize the impact on customers that depend on the existence of the 4-letter message prefixes.

The following is a sample of an entry in a trace.log file using a default message ID. Note that the message ID is PMON0001A

```
[1/26/05 10:17:12:529 EST] 0000000a PMIImp1      A  PMON0001A: PMI is enabled
```

A sample of the same entry using a new message ID follows. Note that the message ID is CWPMI0001A. All new WebSphere Application Server message IDs begin with 'CW'.

```
[1/26/05 10:17:12:529 EST] 0000000a PMIImp1      A  CWPMI0001A: PMI is enabled.
```

About this task

If you are using a logging tool that uses the standardized format, you might want to change the default configuration settings to format the logging output appropriately. You will need to change the configuration for each Java virtual machine (JVM) in the cell if you want the output formatting to be the same across application servers.

Procedure

- To configure logging files so that they use the newer, 5-letter error message prefixes for each process, use the following commands with the `wsadmin` utility:

- Using Jacl:

```
set cfgJvmList [$AdminConfig list JavaVirtualMachine]
set cfgJvm [lindex $cfgJvmList JavaVirtualMachine]
$AdminConfig create Property $cfgJvm {{name com.ibm.websphere.logging.messageId.version} {value 6} {required false}}
$AdminConfig save
```

- Using Jython:


```
ls = java.lang.System.getProperty("line.separator")
cfgJvmList = AdminConfig.list("JavaVirtualMachine").split(ls)
print cfgJvmList
cfgJvm = cfgJvmList[JavaVirtualMachine]
AdminConfig.create('Property', cfgJvm, [['name', 'com.ibm.websphere.logging.messageId.version'], ['value', '6'],
['required', 'false']])
AdminConfig.save()
```

Where `JavaVirtualMachine` is the number of the process that you want to use.

When you specify the process, the first process listed is zero (0), the second process is one (1), and so on. Make the changes for each JVM in the cell for consistent output formatting.

Important: Restart the application server for the changes to take effect.

- To change the configuration so that the log files contain the newer, 5–letter message prefixes in the `startServer.log` or `stopServer.log` files, modify the `startServer` and `stopServer` scripts in the `install_root/bin` directory.

 Within these scripts, append the following code to the end of the existing `D_ARGS` parameter:

```
$DEBUG -Dcom.ibm.websphere.logging.messageId.version=6
```

Results

Message IDs written to log files will now be compliant with the new standard.

Converting log files to use IBM unique Message IDs

The `convertlog` command creates a new log file with either new or old message IDs substituted in place of the message IDs in the source file.

Before you begin

Note: Prior to Version 6.x, components were assigned message IDs that are not necessarily unique across IBM software products. In Version 6.0, a system property was provided to map the message IDs in output logs to a set of IBM unique message IDs (all WebSphere Application Server message IDs now start with CW) that do not conflict with other IBM software products. The default runtime behavior still uses the old message IDs.

About this task

To facilitate the migration of logging tools that are reliant on the old message IDs, the `convertlog` command is provided to convert the message IDs of log entries from the old standard to the new standard, or the new standard back to the old. By default, the software is configured to use the old message IDs when logging, but you can change the default output with the `com.ibm.websphere.logging.messageId.version` system property. Read “Changing the message IDs used in log files” on page 92 for more information.

Procedure

Use the `convertlog` command to convert the log output:

```
convertlog <source file name> <destination file name> [options]
options: -newMessageFormat convert message IDs to CCCCnNnnS format
         (cannot be used with -m5)
         -oldMessageFormat convert message IDs to CCCCnNnnS format
         (cannot be used with -m6)
```

Results

After using the `convertlog` command you have a new file with message IDs in the chosen format.

convertlog command

The `convertlog` command is used to convert the message IDs in log entries from the old standard to the new standard, or the new standard back to the old.

Previous versions of WebSphere Application Server used message IDs that are deprecated in WebSphere Application Server Version 8.5. To facilitate the migration of tools based on the old message IDs, the `convertlog` command is implemented to translate log files from one message ID standard to the other.

Use the `convertlog` command as follows:

```

convertlog <source file name> <destination file name> [options]
  options: -newMessageFormat convert message IDs to CCCCnnnnS format
           (cannot be used with -m5)
           -oldMessageFormat convert message IDs to CCCCnnnnS format
           (cannot be used with -m6)

```

MessageConverter class

The `com.ibm.websphere.logging.MessageConverter` class provides a method to convert a message ID at the front of a `String` into either a new message ID or an old message ID. The direction of the conversion is controlled with the `conversionType` argument.

Use the `MessageConverter` class with log analysis tools to convert message IDs from earlier versions of WebSphere Application Server into the corresponding message IDs that are used in later releases, or to revert message IDs to an earlier format.

Method:

```
public static java.lang.String convert(java.lang.String in, short conversionType)
```

Parameters:

Use the following parameters with the `MessageConverter` class:

Parameter Name	Description
<i>in</i>	The message to convert. The method assumes the message ID is the first part of the supplied message with no leading white space.
<i>conversionType</i>	CONVERSION_TYPE_WASV5_TO_WASV6 CONVERSION_TYPE_WASV6_TO_WASV5

HTTP error, FRCA, and NCSA access log settings

Use this page to configure the global HTTP error log, and National Center for Supercomputing Applications (NCSA) access log settings for an HTTP inbound channel. If you are running the product on z/OS, you can also use this page to configure the global Fast Response Cache Accelerator (FRCA) log settings for an HTTP inbound channel. A FRCA log is a specialized form of a NCSA log and can only be created in a z/OS environment.

To view this administrative console page, click **Servers > Server Types > WebSphere application servers > *server_name***. Under **Troubleshooting**, click **NCSA access and HTTP error logging**. This console page has separate sections for each type of logging. The FRCA logging section only appears if you are running the product on z/OS.

The HTTP error log contains a record of HTTP processing errors that occur. The level of error logging that occurs is dependent on the value that is selected for the Error log level field.

The NCSA access log contains a record of all inbound client requests that the HTTP transport channel handles. All of the messages that are contained in a NCSA access log are in NCSA format.

Configuring and enabling the logging is a two step process. After you use this page to configure the logging, you must explicitly enable each type of logging for the appropriate HTTP channels. To view the settings page for an HTTP channels, click **Servers > Server Types > WebSphere application servers > *server* > Web Container Settings > Web container transport chains > Chain > HTTP inbound channel**.

gotcha: The settings for any of these logs can also be modified on the settings page for a specific HTTP inbound channel. Any changes that you make on the HTTP inbound channel settings page only apply to that specific inbound channel and override any global configuration settings that you specify on this page.

Enable logging service at server start-up

Select this option if you want any of the following logging to start when the server starts:

- NCSA access logging
- HTTP error logging

NCSA access logging

Enable access logging

When selected, a record of inbound client requests that the HTTP transport channel handles is kept in the NCSA access log.

Access log file path

Specifies the directory path and name of the NCSA access log. Standard variable substitutions, such as `$(SERVER_LOG_ROOT)`, can be used when specifying the directory path.

Access log maximum size

Specifies the maximum size, in megabytes, of the NCSA access log. When the content of the NCSA access log reaches the specified maximum size limit, a `<logname>.<timestamp>.log` archive file is created. The current content of the NCSA access log is then copied to this archive log.

An example of a file name for this archive log follows::

```
nca_access_11_09_20_16.15.04.log
```

The next time the content in the NCSA access log reaches the specified maximum log size, the content of the NCSA access log is again copied to the `<logname>.<timestamp>.log` archive file. The copy process overwrites the current content of the archive file with the most current content of the NCSA access log.

NOTE: When there are multiple archive logs, as determined by the setting of the "Maximum number of historical files", the oldest archive log is the one overwritten.

Maximum number of historical files

Specifies the maximum number of historical versions of the NCSA access log file that are kept for future reference.

You can use the `EnableBuildBackupList` HTTP transport custom property to enable the HTTP channel to scan for the history files in the access and error logs directory, and rolling these files over with any newer log files created. See the topic *HTTP Transport channel custom properties* for a description of how to specify this custom property.

NCSA access log format

Specifies which NCSA format is used when logging client access information. If you select `Common`, the log entries contain the requested resource and a few other pieces of information, but does not contain referral, user agent, and cookie information. If you select `Combined`, referral, user agent, and cookie information is included.

Entries in the NCSA access log contain a local time stamp.

You can use the HTTP transport channel custom property `accessLogFormat` to customize the format of the NCSA access log for a specific HTTP transport channel. See the topic *HTTP transport channel custom properties* for a description of how to use this custom property.

Error logging

Enable error logging

When selected, HTTP errors that occur while the HTTP channel processes client requests are recorded in the HTTP error log.

Log file path

Specifies the directory path and the name of the HTTP error log. Standard variable substitutions, such as \$(SERVER_LOG_ROOT), can be used when specifying the directory path.

Error log maximum size

Specifies the maximum size, in megabytes, of the HTTP error log. When the content of the HTTP error log reaches the specified maximum size limit, a *<logname>.<timestamp>.log* archive file is created. The current content of the HTTP error log is then copied to this archive log.

An example of a file name for this archive log follows::

```
http_access_11_09_20_16.15.04.log
```

The next time the content in the HTTP error log reaches the specified maximum log size, the content of the HTTP error log is again copied to the *<logname>.<timestamp>.log* archive file. The copy process overwrites the current content of the archive file with the most current content of the HTTP error log.

NOTE: When there are multiple archive logs, as determined by the setting of the "Maximum number of historical files", the oldest archive log is the one overwritten.

Maximum number of historical files

Specifies the maximum number of historical versions of the Error log file that are kept for future reference.

You can use the EnableBuildBackupList HTTP transport custom property to enable the HTTP channel to scan for the history files in the access and error logs directory, and rolling these files over with any newer log files created. See the topic *HTTP Transport channel custom properties* for a description of how to specify this custom property.

Error log level

Specifies the type of error messages that are included in the HTTP error log.

You can select:

Critical

Only critical failures that stop the Application Server from functioning properly are logged.

Error The errors that occur in response to clients are logged. These errors require Application Server administrator intervention if they result from server configuration settings.

Warning

Information on general errors, such as socket exceptions that occur while handling client requests, are logged. These errors do not typically require Application Server administrator intervention.

Information

The status of the various tasks that are performed while handling client requests is logged.

Debug

More verbose task status information is logged. This level of logging is not intended to replace RAS logging for debugging problems, but does provide a steady status report on the progress of individual client requests. If this level of logging is selected, you must specify a large enough log file size in the Error log maximum size field to contain all of the information that is logged.

Chapter 6. Using High Performance Extensible Logging to troubleshoot applications

You can use High Performance Extensible Logging (HPEL) to help diagnose problems in WebSphere Application Server.

About this task

Administrators using WebSphere Application Server need to use log and trace files to determine whether their applications and the server are running correctly.

Logs typically contain information that is of interest to administrators and must be looked at periodically to ensure there are no unexpected errors or warnings.

Trace typically contains information that is useful for debugging application or server problems and can help identify specific problems with individual components.

Procedure


1. Enable HPEL if you have not done so already. Read about changing from basic mode to HPEL logging and tracing for more information.
2. Configure the HPEL facility. For example, configure HPEL to store your logs and trace in appropriate directories, and specify how long you want log and trace content to be retained before being deleted. Read about HPEL to learn about the log and trace framework overall, and read about configuring HPEL for more configuration information.
3. Restart the application server after making static configuration changes.

High Performance Extensible Logging (HPEL)

High Performance Extensible Logging (HPEL) is a log and trace facility that is provided as a part of WebSphere Application Server.

Overview

Note: The basic log and trace facility is enabled by default. To use HPEL you must enable it.

 HPEL provides a convenient mechanism for storing and accessing log, trace, System.err, and System.out information produced by the application server or your applications. It is an alternative to the basic log and trace facility, which provided the JVM logs, diagnostic trace, and service log files commonly named SystemOut.log/SystemErr.log, trace.log, and activity.log.

HPEL log and trace storage

HPEL provides a log data repository, a trace data repository, and a text log file. See the following figure to understand how applications and the application server store log and trace information.

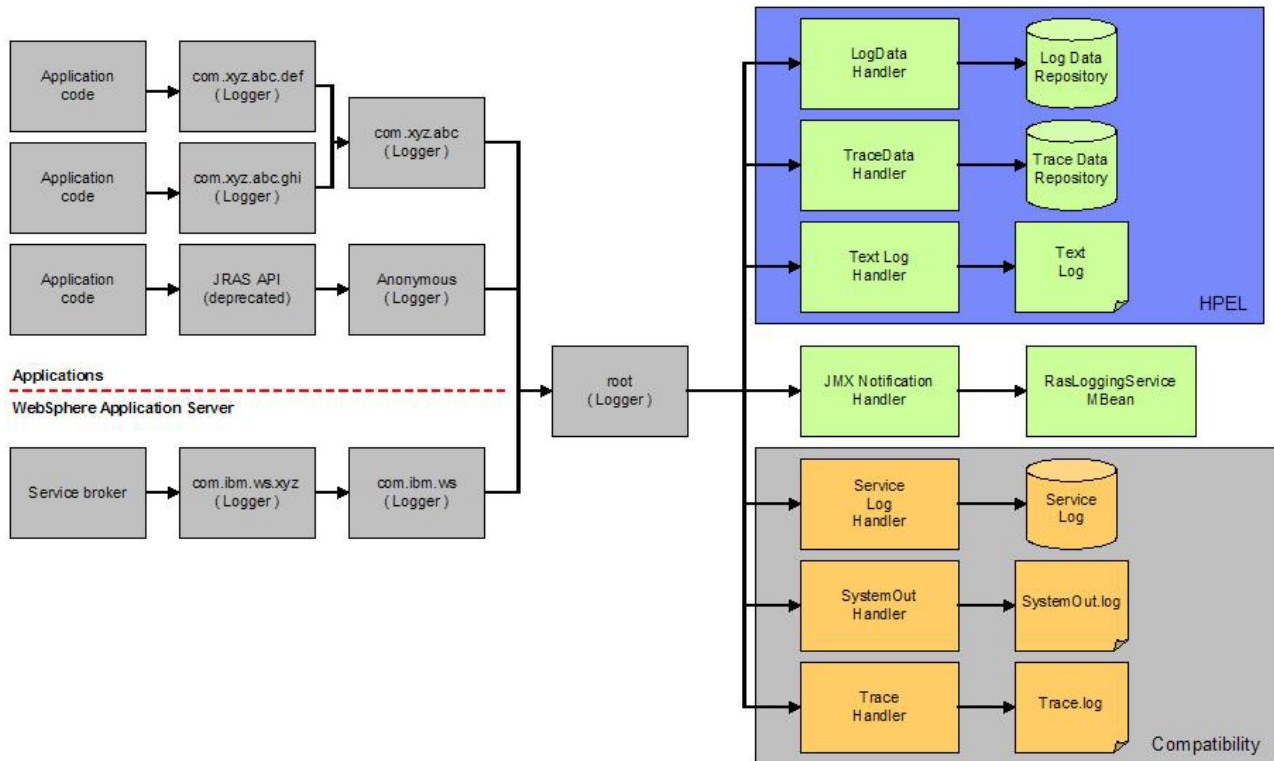


Figure 5. Log and trace storage for HPEL and basic logging.

HPEL log data repository

The log data repository is a storage facility for log records. Log data is typically intended to be reviewed by administrators. This includes any information applications or the server write to System.out, System.err, or java.util.logging at level Detail or higher (including Detail, Config, Info, Audit, Warning, Severe, Fatal, and any custom levels at level Detail or higher).

HPEL trace data repository

The trace data repository is a storage facility for trace records. Trace data is typically intended for use by application programmers or by the WebSphere Application Server support team. This includes any information applications or the server write to java.util.logging at levels below level Detail (including Fine, Finer, Finest, and any custom levels below level Detail).

Note: Log and trace content written to the deprecated JRAS logging API is also included in the log and trace data repositories. Some logging APIs, such as Jakarta Commons Logging can also be configured to route their log and trace data to java.util.logging, and would have their output stored in the log data or trace data repository as well.

HPEL text log

The text log file is a plain text file for log and trace records. The text log file is provided for convenience, primarily so that log content can be read without having to run the LogViewer command-line tool to convert the log data repository content to plain text.

The text log file does not contain any content that is not also stored in either the log data repository or trace data repository. You can disable the text log to enhance server performance. The text log can be configured to record trace content for debugging convenience.

Note: Writing trace to the text log is expensive from a performance perspective.


Log and trace performance

IBM i HPEL has been designed and tested to significantly outperform the existing basic log and trace facility. One result is that the application server can run with trace enabled while causing less impact to performance than tracing the same components using basic logging. Another result is that applications that frequently write to the logs might run faster with HPEL. A number of factors contribute to the overall performance of HPEL logging and tracing.

Log and trace events are each stored in only one place

Log events, System.out, and System.err are stored in the log data repository. Trace events are stored in the trace data repository. If the text log file is disabled, HPEL might only write log and trace content to these repositories. Storing each type of event in one place ensures that performance is not wasted on redundant data storage.

Log events, and optionally trace events, are written to the text log file when it is enabled. Since this data is always also stored in the log data and trace data repositories, the text log file content is redundant. The text log is convenient for users who do not want to run the LogViewer command-line tool to see their logs and trace; but you can disable the text log if this convenience is not needed.



Log and trace repositories are not shared across processes

Synchronizing activities between processes causes a degradation in performance to all processes involved. With HPEL, each server process has its own log data repository, trace data repository, and text log file. Since these files are not shared across processes, the server runtime environment does not need to synchronize with other processes when writing to these destinations.

Data is not formatted unless it is needed

Formatting data for a user to read uses processor time. Rather than format log event and trace event data at run time, HPEL log and trace data is stored more rapidly in a proprietary binary representation. This improves the performance of the log and trace facility. By deferring log and trace formatting until the LogViewer is run, sections of the log or trace that are never viewed are never formatted.

You can enable the text log file, which stores the log data and trace data in an already readable text format.

Note: Disable the text log when performance of your server is a key concern, or if the text log is not wanted.

Log and trace data is buffered before being written to disk

Writing large blocks of data to a disk is more efficient than writing the same amount of data in small blocks. HPEL provides buffer log and trace data before writing it to disk. By default, log and trace data is stored in an 8 KB buffer before being written to disk. If the buffer is filled within 10 seconds, the buffer is written to disk. If the buffer is not filled within that time it is automatically written to disk to ensure that the logs have the most current information.


Note: The size of the buffer can be controlled by setting the HPEL.BUFFER.SIZE system property. The frequency that the buffer writes to disk can be controlled by setting the HPEL.FLUSH.PERIOD.SECONDS system property.

Administration of log and trace

HPEL has been designed to be easy to configure and understand. For example, administrators can easily configure how much disk space to dedicate to logs or trace, or how long to retain log and trace records, and leave the management of log and trace content up to the server. As another example all log, trace, System.out, and System.err content can be accessed using one easy-to-use command (LogViewer), avoiding any possible confusion over which file to access for certain content.

Reading from the log data and trace data repositories

The log data and trace data repositories are stored in a WebSphere Application Server proprietary format and cannot be read using text file editors such as Notepad or VI. You can copy the log data and trace data repositories in to a plain text format using the LogViewer command.



HPEL LogViewer command

The HPEL LogViewer is an easy-to-use, command-line tool provided for HPEL users to work with the log data and trace data repositories. The LogViewer provides filtering and formatting options that make finding important content in the log data and trace data repositories easy. For example, a user might filter any errors or warnings, then filter all log and trace entries that occurred within 10 seconds of a key error message on the same thread.

Filtering using log and trace record extension content

Note: You can use the LogViewer command-line tool to filter records based on the content of log and trace record extensions. The application server automatically creates an appName extension for each log and trace record related to a Java Platform, Enterprise Edition (Java EE) application, indicating the name of that application. The application server also automatically creates a requestID extension for each log and trace record created during the processing of certain types of requests (for example HTTP or JMS requests), indicating the unique ID of that request.

The requestID extension is added only to log and trace records when Cross-Component Trace is enabled. HPEL also provides the ability for developers to add custom extensions to log and trace records using a log record context API (com.ibm.websphere.logging.hpel.LogRecordContext).

HPEL in the administrative console

The administrative console contains pages that enable HPEL administrators to:

- Configure the HPEL log data repository.

- Configure the HPEL trace data repository.
- Configure the HPEL text log file.
- View the contents of the HPEL log and trace data repositories.
- View and set the log detail levels for logging and tracing.

To use these capabilities, in the administrative console, click **Troubleshooting > Logs and Trace** link.


Development resources

HPEL has been designed to make working with log and trace content more flexible and effective than the basic logging facility. Log and trace content can be easily filtered to show only the records that are of interest. You can use the command line (see the description of the HPEL LogViewer command), or developers can create powerful log handling programs using the HPEL API.

Scripts and Java programs read from the log data and trace data repositories

Developers and scripters have a number of options for how to read the log data and trace data repositories:

- Locally or remotely from a `wsadmin` script, using the HPELControlService JMX MBean
- Locally or remotely from a Java program, using the HPELControlService JMX MBean
- Locally from a Java program, using the `com.ibm.websphere.logging.hpel` API



HPEL-related JMX MBeans

A MBean interface has been provided to make it easy to access HPEL repository content remotely. For example, a developer might write a JMX client program to read log content from across their WebSphere Application Server cell. This interface is part of the HPELControlService MBean. Refer to the MBean interface documentation for details on the HPEL remote log reading interface.

Table 18. JMX MBeans related to HPEL. The JMX MBeans are related to the operation of HPEL.

JMX MBean	Description
HPELControlService	Provides operations related to configuring the log or trace detail level of the server, viewing the log component registry, and querying the log and trace repositories
HPELLogDataService	Provides operations related to configuring the log data repository of the server
HPELTraceDataService	Provides operations related to configuring the trace data repository of the server
HPELTextLogService	Provides operations related to configuring the text log file of the server
RasLoggingService	Only used for JMX Notification of log events

When using HPEL for log and trace rather than basic logging, the log and trace JMX MBean, TraceService, is not used.

HPEL API

An API has been provided to make it easy for developers to develop tools to consume content from the HPEL log and trace repositories. For example, a developer might write a Java program to search the log and trace content to find any messages with message IDs that match a known list of important message IDs. This API is in the `com.ibm.websphere.logging.hpel` package. Refer to the API documentation for details on the HPEL log reading API.

Log and trace record extensibility

Note: Developers can use HPEL to add custom extensions to log and trace records through a log record context API (`com.ibm.websphere.logging.hpel.LogRecordContext`). When HPEL stores log and trace records, it includes any extensions present in the log record context on the same thread. For example, a developer might write a servlet filter to add important HTTP request parameters to the log record context. While that servlet runs, HPEL adds those extensions to any log and trace records created on the same thread.

As with other log and trace record fields, developers can access the record extensions using the HPEL API. This is useful when writing tools to read from log and trace repositories. Developers can also make use of the log record context API to access extensions in custom log handlers, filters, and formatters at run time.

Basic mode and HPEL mode

Two modes of logging and tracing exist in the product, which are basic mode and High Performance Extensible Logging (HPEL) mode. Use this topic to understand the differences between these modes.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Basic mode, the default mode, is the existing log and trace framework from prior releases of WebSphere Application Server. Any existing scripts and tools you have that use logs and trace might need minor modifications to work with HPEL mode.

To take advantage of the new log and trace framework, HPEL mode must be enabled. Once HPEL mode is enabled the JVM logs (typically `SystemOut.log` and `SystemErr.log`), the trace log (typically `trace.log`), and the service log (typically `activity.log`) are no longer written to. Instead, log and trace content is written to a log data or trace data repository in a proprietary binary format and, if configured, to a text log file. Disabling the writing of the text log file results in the largest possible performance benefit of HPEL. A log viewing tool, LogViewer, is provided to allow for viewing, filtering, monitoring, and formatting the log and trace data in the repositories.

The following figure shows the files used by the basic mode and HPEL mode log and trace facilities. When enabled, the HPEL text log file stores content from Java trace (optional), Java logs, `System.out`, and `System.err`. You can disable the HPEL text log in cases where it is not needed as indicated by the dotted lines.

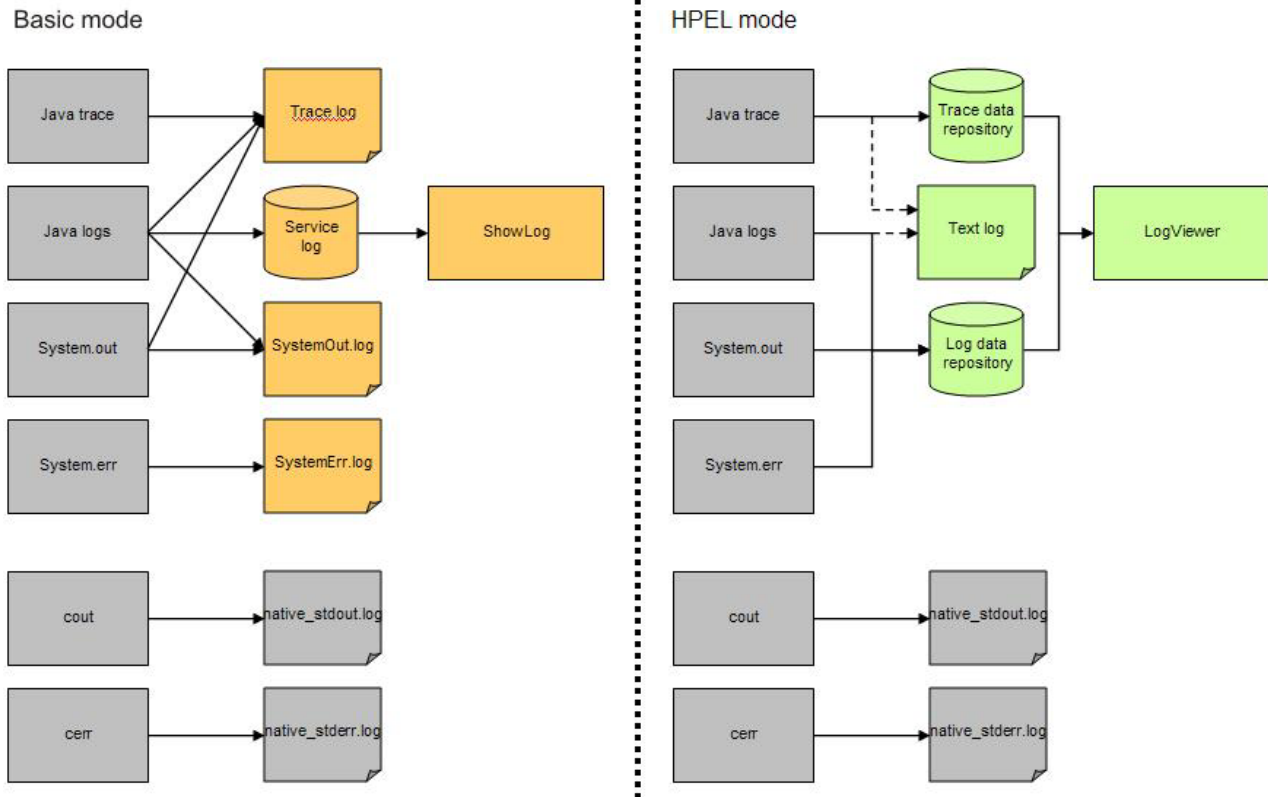


Figure 6. Log and trace files

Table 19. Basic mode and HPEL mode files. This table shows where each of the application and application server log, trace, JVM log, and native JVM log content is routed and how to view that content.

Source	Basic mode files	HPEL mode files	How to view the HPEL files
System.out	SystemOut.log trace.log (when trace enabled)	logdata/*.wbl TextLog_<timestamp>.log (when text log enabled)	logdata - use LogViewer, with optional filtering, to render log data repository as readable text, then use any text editor TextLog - any text editor
System.err	SystemErr.log	Same as System.out	Same as System.out
java.util.logging (levels DETAIL and above)	SystemOut.log activity.log trace.log (when trace enabled)	Same as System.out	Same as System.out
java.util.logging (levels below DETAIL)	trace.log	tracedata/*.wbl TextLog_<timestamp>.log (when text log enabled)	tracedata - use LogViewer, with optional filtering, to render trace data repository as readable text, then use any text editor TextLog - any text editor
native cout	native_stdout.log	native_stdout.log	Any text editor
native cerr	native_stderr.log	native_stderr.log	Any text editor

The following table describes the MBean services:

Table 20. Basic mode and HPEL mode MBeans. This table shows the JMX MBeans that are associated with the basic mode and HPEL mode log and trace frameworks.

Basic mode MBeans	HPEL mode MBeans	HPEL MBean Descriptions
RasLoggingServiceMBean	RasLoggingServiceMBean	Provides JMX Notification
	HPELLogDataServiceMBean	Configures the log repository such as location, retention policy, out of space behavior, buffering, and file switching
	HPELTraceDataServiceMBean	Configures the trace repository such as location, retention policy, out of space behavior, buffering, and file switching
	HPELTextLogServiceMBean	Configures the text log such as location, retention policy, out of space behavior, buffering, file switching, and SystemErr or SystemOut format
TraceServiceMBean	HPELControlServiceMBean	Configures trace specification levels, and provides access to log and trace repository content

Changing from basic mode to HPEL logging and tracing

The basic mode log and trace framework is enabled by default when you set up a new application server. Use this topic to switch to the High Performance Extensible Logging (HPEL) log and trace framework.

Before you begin

Before beginning this task, read about the differences between HPEL mode and basic mode. Be aware of changes you might need to make to any tools and scripts you have that use the basic mode log and trace files.

About this task

HPEL provides faster log and trace handling capabilities and more flexible ways to use log and trace content than the basic mode. You can switch to HPEL mode using the administrative console, or using wsadmin scripting.

Procedure

- Use the administrative console to switch to HPEL.
 1. Log on to the administrative console.
 2. If using an admin agent topology select a node that you want to manage and navigate to it.
 3. From the navigation section in the console choose **Troubleshooting > Logs and trace** .
 4. Select the server that you want to switch to HPEL.
 5. Click **Switch to HPEL Mode** .
 6. Save the changes.
- Use wsadmin scripting to switch to HPEL. Complete these steps to modify the server configuration.
 1. Start wsadmin. In this case, you can connect wsadmin to a running server or access the configuration data for a stopped server. Read about starting the wsadmin scripting client for more information.
 2. Get a reference to the HighPerformanceExtensibleLogging configuration object.

Using Jython:

```
HPELService = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/HighPerformanceExtensibleLogging:/")
```

Table 21. AdminConfig command description. The table lists AdminConfig command and description.

Command	Description
<i>myCell</i>	The name of the cell
<i>myNode</i>	The host name of the node
<i>myServer</i>	The name of the server

3. Set the HighPerformanceExtensibleLogging enable attribute to true.

Using Jython:

```
AdminConfig.modify(HPELService, "[[enable true]]")
```

4. Get a reference to the RASLoggingService object.

Using Jython:

```
RASLogging = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/RASLoggingService:/")
```

Table 22. AdminConfig command description. The table lists AdminConfig command and description.

Command	Description
<i>myCell</i>	The name of the cell
<i>myNode</i>	The host name of the node
<i>myServer</i>	The name of the server

5. Set the RASLoggingService enable attribute to false.

Using Jython:

```
AdminConfig.modify(RASLogging, "[[enable false]]")
```

6. Save the configuration.

Using Jython:

```
AdminConfig.save()
```

Results

The server is now configured to use HPEL when you restart.

What to do next

Configure HPEL to meet your needs. If you have any tools or scripts that were using the basic mode log and trace files, you might need to modify them to continue working with HPEL.

Changing from HPEL to basic mode logging and tracing

Use this topic if you need to switch from HPEL to the basic mode log and trace framework. The basic mode log and trace framework is enabled by default when you set up a new application server.

Before you begin

Before beginning this task, read about the differences between HPEL mode and basic mode. Be aware of changes you might need to make to any tools and scripts you have that use HPEL files or commands.

About this task

You can switch to basic mode using the administrative console, or using wsadmin scripting. HPEL provides faster log and trace handling capabilities and more flexible ways to exploit log and trace content than the basic mode.

Procedure

- Use the administrative console to switch to basic mode.
 1. Log on to the administrative console.
 2. If using an administrative agent topology, select a node that you want to manage and navigate to it.
 3. From the navigation section in the console click **Troubleshooting > Logs and trace**.
 4. Select the server that you want to switch to basic mode.
 5. Click **Change log and trace mode**.
 6. Click **Switch back to basic mode**.
 7. Save the changes.
- Use wsadmin scripting to switch to basic mode. These steps modify the server configuration. The server does not need to be running to perform these steps.
 1. Start wsadmin. In this case, wsadmin can be connected to a running server or accessing the configuration data for a stopped server. Read about starting the wsadmin scripting client for more information.
 2. Get a reference to the HighPerformanceExtensibleLogging configuration object.

Using Jython:

```
HPPELService = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/HighPerformanceExtensibleLogging:/")
```

Table 23. AdminConfig command description. The table lists AdminConfig command and description.

Command	Description
<i>myCell</i>	The name of the cell
<i>myNode</i>	The host name of the node
<i>myServer</i>	The name of the server

3. Set the HighPerformanceExtensibleLogging enable attribute to false.

Using Jython:

```
AdminConfig.modify(HPPELService, "[[enable false]]")
```

4. Get a reference to the RASLoggingService object.

Using Jython:

```
RASLogging = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/RASLoggingService:/")
```

Table 24. AdminConfig command description. The table lists AdminConfig command and description.

Command	Description
<i>myCell</i>	The name of the cell
<i>myNode</i>	The host name of the node
<i>myServer</i>	The name of the server

5. Set the RASLoggingService enable attribute to true.

Using Jython:

```
AdminConfig.modify(RASLogging, "[[enable true]]")
```

6. Save the configuration.

Using Jython:

```
AdminConfig.save()
```

Results

The server is now configured to use the basic mode log and trace framework when you restart.

Determining which of basic mode and HPEL mode is enabled

WebSphere Application Server offers both a High Performance Extensible Logging (HPEL) log and trace framework, and a basic log and trace framework. There are a number of ways to determine which of the two frameworks is enabled.

About this task

In many circumstances it might be important to know whether the HPEL or basic log and trace framework is enabled, for example, when writing a script whose purpose is to read from any log files of the server.

Since the configuration files of a server can differ from the running state of a server (for example when configuration changes have been made but the server has not been restarted), steps are provided for determining the log and trace mode in various ways.

Procedure

- Use wsadmin to determine the log and trace mode that a running server uses.
 1. Start wsadmin. In this case, wsadmin must be connected to a running server, for example through the SOAP port. Read about starting the wsadmin scripting client for more information.
 2. Determine whether the HPELControlService object is available. If the HPELControlService is present it can be concluded that the server is running with the HPEL log and trace framework.

Using Jython:

```
HPELMBean = AdminControl.queryNames('cell=myCell,node=myNode,
type=HPELControlService,process=myServer,*')
if (HPELMBean == ''):
    print "HPEL is not enabled"
else:
    print "HPEL is enabled"
```

Table 25. AdminControl command description. The table lists AdminControl command and their description.

Command	Description
<i>myNode</i>	The host name of the node
<i>myServer</i>	The name of the server

- Use wsadmin to determine the log and trace mode that a servers configuration specifies.
 1. Start wsadmin. In this case, wsadmin can be connected to a running server or accessing the configuration data for a stopped server. Read about starting the wsadmin scripting client for more information.
 2. Determine whether the RASLoggingService configuration object is enabled. If the RASLoggingService config object is enabled then it can be concluded that the server is configured to run with the basic log and trace framework. Otherwise, if the HighPerformanceExtensibleLogging config object is enabled it can be concluded that the server is configured to run with the HPEL log and trace framework.

Note: If both the RASLoggingService config object and the HighPerformanceExtensibleLogging config object are enabled it can be concluded that the server is configured to run with the basic log and trace framework.

Using Jython:

```
RASLogging = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/RASLoggingService:/")
basicEnabled = AdminConfig.showAttribute(RASLogging, "enable")
if (basicEnabled == "true"):
    print "Basic mode logging in effect"
else:
    HPELsvc = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/HighPerformanceExtensibleLogging:/")
    HpelEnabled = AdminConfig.showAttribute(HPELsvc, "enable")
```

```

if (HpelEnabled == "true"):
    print "HPEL is enabled"
else:
    print "No logging is enabled"

```

Table 26. AdminControl command description. The table lists AdminControl command and their description.

Command	Description
<i>myCell</i>	The name of the cell
<i>myNode</i>	The host name of the node
<i>myServer</i>	The name of the server

- Use the administrative console to determine the log and trace mode that a running server uses.
 1. Log into the administrative console.
 2. Click **Troubleshooting > Logs and Trace > *myServer*** (where *myServer* is the name of the server you are interested in)
 3. Find the **Switch to HPEL Mode** button. If this button is available, the server is using the basic log and trace framework. Otherwise, the server is using HPEL.

Results

For any method selected, the result is that you now know whether a server is configured to use the HPEL or basic mode log and trace framework.

Configuring HPEL with wsadmin scripting

You can configure the High Performance Extensible Logging (HPEL) log and trace framework using wsadmin scripting. Use the examples in this topic as a guide to build your own wsadmin scripts.

About this task

HPEL provides faster log and trace handling capabilities and more flexible ways to use log and trace content than the basic mode. You can configure the HPEL mode using the administrative console, or using wsadmin scripting. The examples in this topic show how to configure HPEL using wsadmin. If you complete this task using the deployment manager, then you might need to synchronize the node agent on the target node and restart the server before configuration changes take effect.

Table 27. Variable Names. The table applies to all examples in this topic. All examples use the Jython scripting language.

Variable	Description
<i>myCell</i>	The name of the cell
<i>myNode</i>	The host name of the node
<i>myServer</i>	The name of the server

Procedure

- Use the AdminConfig object to configure HPEL.

Changes you make using the AdminConfig object take effect the next time you start the server.

1. Change the trace specification.

The following example shows how to change the trace specification to

```
*=info:com.ibm.ws.classloader.*=all
```

```

HPELService = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/
HighPerformanceExtensibleLogging:/")
AdminConfig.modify(HPELService, "[[startupTraceSpec *=info:com.ibm.ws.classloader.*=all]]")
AdminConfig.save()

```

2. Change the size of the log repository.

The following example shows how to set HPEL to automatically delete the oldest log content from the log repository when the repository size approaches 65 MB. Specify HPELTrace or HPELTextLog instead of HPELLog to change the setting for the HPEL trace repository or HPEL text log.

```
HPELService = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/  
HighPerformanceExtensibleLogging:/")  
HPELLog = AdminConfig.list("HPELLog", HPELService)  
AdminConfig.modify(HPELLog, "[[purgeMaxSize 65]]")  
AdminConfig.save()
```

3. Change the log repository location.

The following example shows how to change the HPEL log repository directory name to /tmp/myDirectory. Specify HPELTrace or HPELTextLog instead of HPELLog to change the setting for the HPEL trace repository or HPEL text log.

```
HPELService = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/  
HighPerformanceExtensibleLogging:/")  
HPELLog = AdminConfig.list("HPELLog", HPELService)  
AdminConfig.modify(HPELLog, "[[dataDirectory /tmp/myDirectory]]")  
AdminConfig.save()
```

4. Disable log record buffering.

The following example shows how to change the HPEL log repository to not use log record buffering. Specify HPELTrace or HPELTextLog instead of HPELLog to change the setting for the HPEL trace repository or HPEL text log.

```
HPELService = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/  
HighPerformanceExtensibleLogging:/")  
HPELLog = AdminConfig.list("HPELLog", HPELService)  
AdminConfig.modify(HPELLog, "[[bufferingEnabled false]]")  
AdminConfig.save()
```

Note: Enable log record buffering in almost all cases. Only disable log record buffering when your server is failing unexpectedly and cannot write buffered content to disk before stopping.

5. Start writing to a new log file each day at a specified time.

The following example shows how to enable the HPEL log repository to start a new log file each day at 3pm. Specify HPELTrace or HPELTextLog instead of HPELLog to change the setting for the HPEL trace repository or HPEL text log.

```
HPELService = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/  
HighPerformanceExtensibleLogging:/")  
HPELLog = AdminConfig.list("HPELLog", HPELService)  
AdminConfig.modify(HPELLog, "[[fileSwitchTime 15]]")  
AdminConfig.modify(HPELLog, "[[fileSwitchEnabled true]]")  
AdminConfig.save()
```

6. Change the out of space action for the log repository.

The following example shows how to change the out of space action for the HPEL log repository. Specify HPELTrace or HPELTextLog instead of HPELLog to change the setting for the HPEL trace repository or HPEL text log.

```
HPELService = AdminConfig.getid("/Cell:myCell/Node:myNode/Server:myServer/  
HighPerformanceExtensibleLogging:/")  
HPELLog = AdminConfig.list("HPELLog", HPELService)  
AdminConfig.modify(HPELLog, "[[outOfSpaceAction PurgeOld]]")  
AdminConfig.save()
```

- Use the AdminControl object to configure HPEL. Changes you make using the AdminControl object take effect immediately.

1. Change the trace specification.

The following example shows how to change the trace specification to
=info:com.ibm.ws.classloader.=all

```

HPELControlMBean = AdminControl.queryNames('cell=myCell,node=myNode,
type=HPELControlService,process=myServer,*')
AdminControl.setAttribute(HPELControlMBean, "traceSpecification",
"*=info:com.ibm.ws.classloader.*=all")

```

2. Change the size of the log repository.

The following example shows how to set HPEL to automatically delete the oldest log content from the log repository when the repository size approaches 65 MB. Specify HPELTraceDataService or HPELTextLogService instead of HPELLogDataService to change the setting for the HPEL trace repository or HPEL text log.

```

HPELLogDataMBean = AdminControl.queryNames('cell=myCell,
node=myNode,type=HPELLogDataService,process=myServer,*')
AdminControl.setAttribute(HPELLogDataMBean, "purgeMaxSize", "65")

```

3. Change the log repository location.

The following example shows how to change the HPEL log repository directory name to /tmp/myDirectory. Specify HPELTraceDataService or HPELTextLogService instead of HPELLogDataService to change the setting for the HPEL trace repository or HPEL text log.

```

HPELLogDataMBean = AdminControl.queryNames('cell=myCell,
node=myNode,type=HPELLogDataService,process=myServer,*')
AdminControl.setAttribute(HPELLogDataMBean, "dataDirectory", "/tmp/myDirectory")

```

4. Disable log record buffering.

The following example shows how to change the HPEL log repository to not use log record buffering. Specify HPELTraceDataService or HPELTextLogService instead of HPELLogDataService to change the setting for the HPEL trace repository or HPEL text log.

```

HPELLogDataMBean = AdminControl.queryNames('cell=myCell,node=myNode,
type=HPELLogDataService,process=myServer,*')
AdminControl.setAttribute(HPELLogDataMBean, "bufferingEnabled", "false")

```

Note: Enable log record buffering in almost all cases. Only disable log record buffering when your server is failing unexpectedly and cannot write buffered content to disk before stopping.

5. Start writing to a new log file each day at a specified time.

The following example shows how to enable the HPEL log repository to start a new log file each day at 3pm. Specify HPELTrace or HPELTextLog instead of HPELLog to change the setting for the HPEL trace repository or HPEL text log.

```

HPELLogDataMBean = AdminControl.queryNames('cell=myCell,node=myNode,
type=HPELLogDataService,process=myServer,*')
AdminControl.setAttribute(HPELLogDataMBean, "fileSwitchTime", "15")
AdminControl.setAttribute(HPELLogDataMBean, "fileSwitchEnabled", "true")

```

6. Change the out of space action for the log repository.

The following example shows how to change the out of space action for the HPEL log repository. Specify HPELTraceDataService or HPELTextLogService instead of HPELLogDataService to change the setting for the HPEL trace repository or HPEL text log.

```

HPELLogDataMBean = AdminControl.queryNames('cell=myCell,node=myNode,
type=HPELLogDataService,process=myServer,*')
AdminControl.setAttribute(HPELLogDataMBean, "outOfSpaceAction", "PurgeOld")

```

Results

HPEL is now configured. If you made changes with the AdminConfig command, restart the server to make the changes take effect.

Configuring HPEL

HPEL logging and trace settings

Use this page to view and configure High Performance Extensible Logging (HPEL) logging and trace settings for the server.

Note: You can only access this page when the server is configured to use HPEL log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and trace > *server_name***.

Configure HPEL logging

Use this link to configure HPEL log options. If this server is in running state, then important log configuration values are displayed following the link. These values summarize the current runtime values being used by the server.

Configure HPEL trace

Use this link to configure HPEL trace options. If this server is in running state, then important trace configuration values are displayed following the link. These values summarize the current runtime values being used by the server.

Configure HPEL text log

Use this link to configure HPEL text log options. If this server is in running state, then important text log configuration values are displayed following the link. These values summarize the current runtime values being used by the server.

View HPEL logs and trace

Use this page to view log data from the HPEL repository. You can also use this page to filter and search the repository. You can export the customized view or full repository into a text file or into a new HPEL repository.

Change log detail levels

Use this page to enter a log detail level that specifies the components, packages, or groups to trace, and to configure log and trace correlation settings.

Change log and trace mode

Use this link to switch back to basic mode logging, instead of HPEL mode logging currently enabled for this server.

Manage process logs

WebSphere Application Server processes contain two output streams that are accessible to native code, which runs in a particular process. These streams are the **stdout** and **stderr** streams. Native code, including Java virtual machines (JVM), might write data to these process streams. In addition, you can also configure JVM-provided System.out and System.err streams to write their data to these streams.

NCSA access and HTTP error logging

The NCSA access and HTTP error logging page enables you to configure the log settings for your HTTP server.

HPEL log configuration settings

Use this page to configure High Performance Extensible Logging (HPEL) log settings.

Note: You can only access this page when the server is configured to use HPEL log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and trace > *server_name* > Configure HPEL logging** .

Directory path

Specifies the directory to which log files are written. A subdirectory, logdata, is created in this directory, and the log files are written to this location.

Note: This path cannot be shared between servers. The server creates a file with a .owner extension to help detect when two or more servers happen to be trying to use the same path for HPEL output.

Enable log record buffering

Specifies that the logging system avoids writing to disk each time a log record is created. The logging system creates a buffer that can hold a number of log records, and writes the buffered events when the buffer is full. The logging system also writes the buffered events after a few seconds have passed, even if the buffer is not full.

Selecting this setting significantly improves logging performance; however, if the server stops unexpectedly, the contents might not be written to the log repository.

Note: Enable log record buffering in almost all cases. Only disable log record buffering when your server is failing unexpectedly and cannot write buffered content to disk before stopping.

Start new log file daily at <time>

Enables the logging framework to close the log file and start a new file at the specified time of day. Closing the file makes it easy to copy the file to an archive.

Note: If you want to automatically archive your log files, set up your backup program to copy files after the time you configured for new logs to be started. Configure the backup to occur at least 10 minutes after the time configured for new logs to be started to ensure that the server has closed the previous file.

Begin cleanup of oldest records

Specifies the log cleanup settings to be used to automatically purge the oldest log records, or log records that no longer fit in the configured space, from the log repository.

Select **When log size approaches maximum** to configure automatic log file cleanup to begin when the total size of the log repository approaches the configured maximum size.

Select **When oldest records reach age limit** to configure automatic log file cleanup to begin when log content is the age limit specified.

Select **When either age or size restriction is met** to configure automatic log file cleanup to begin when either of the previous conditions is met.

Regardless of the selection chosen, records are deleted from the log repository in the order in which they were written to the log repository.

Log record age limit:

Specifies the lifespan, in hours, that log records can remain in the log repository before the log records can be automatically deleted by the server. When the oldest records in the log repository have existed longer than the age limit specified, then those records are targeted for deletion by the server.

Maximum log size:

Specifies the maximum total size, in megabytes, that the server allows the log repository to reach. When the log repository approaches this size limit, the server deletes the oldest records from the log repository to make space for new log records.

Out of space action

Specifies how the server reacts to an inability to add content to the log repository.

Select **Stop server** to specify that the server stops when the server is unable to write to the log repository.

Select **Purge old records** to specify that the server continues to run, and that the oldest log records are immediately removed when the server is unable to write to the log repository.

Select **Stop logging** to specify that the server continues to run, but that the server cannot continue to write to the log when the server is unable to write to the log repository.

Save runtime changes to configuration as well

Specifies that changes are made to both the dynamic state of the running server, and the server configuration, which takes effect on the next restart. If this check box is not selected, the server does not copy the settings into the server configuration.

HPEL trace configuration settings

Use this page to configure High Performance Extensible Logging (HPEL) trace settings.

Note: You can only access this page when the server is configured to use HPEL log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and trace > *server_name* > Configure HPEL trace**.

Trace to a directory

Specifies that the tracing system writes trace records to the trace directory as they are created by the server.

Enable log record buffering

Specifies that the tracing system avoids writing to disk each time a trace record is created. The tracing system creates a buffer that can hold a number of trace records, and writes the buffered events when the buffer is full. The tracing system also writes the buffered events after a few seconds have passed, even if the buffer is not full.

Selecting this setting significantly improves tracing performance; however, if the server stops unexpectedly, the contents might not be written to the trace repository.

Note: Enable trace record buffering in almost all cases. Only disable trace record buffering when your server is failing unexpectedly and cannot write buffered content to disk before stopping.

Start new log file daily at <time>

Enables the tracing framework to close the trace file and start a new file at the specified time of day. Closing the file makes it easy to copy the file to an archive.

Note: If you want to automatically archive your trace files, set up your backup program to copy files after the time you configured for new files to be started. Configure the backup to occur at least 10 minutes after the time configured for new files to be started to ensure that the server has closed the previous file.

Begin cleanup of oldest records

Specifies the trace cleanup settings to be used to automatically purge the oldest trace records, or trace records that no longer fit in the configured space, from the trace repository.

Select **When log size approaches maximum** to configure automatic trace file cleanup to begin when the total size of the trace repository approaches the configured maximum size.

Select **When oldest records reach age limit** to configure automatic trace file cleanup to begin when trace content is the age limit specified.

Select **When either age or size restriction is met** to configure automatic trace file cleanup to begin when either of the previous conditions is met.

Regardless of the selection chosen, records are deleted from the trace repository in the order in which they were written to the trace repository.

Log record age limit:

Specifies the lifespan, in hours, that trace records can remain in the trace repository before the trace records can be automatically deleted by the server. When the oldest records in the trace repository have existed longer than the age limit specified, then those records are targeted for deletion by the server.

Maximum log size:

Specifies the maximum total size, in megabytes, that the server allows the trace repository to reach. When the trace repository approaches this size limit, the server deletes the oldest records from the trace repository to make space for new trace records.

Out of space action

Specifies how the server reacts to an inability to add content to the trace repository.

Select **Stop server** to specify that the server stops when the server is unable to write to the trace repository.

Select **Purge old records** to specify that the server continues to run, and that the oldest trace records are immediately removed when the server is unable to write to the trace repository.

Select **Stop logging** to specify that the server continues to run, but that the server cannot continue to write to the trace when the server is unable to write to the trace repository.

Trace to a memory buffer

Specifies that the tracing system writes trace records to a memory buffer.

You can write the memory buffer contents to the trace directory from the runtime tab.

Memory buffer size

Specifies the amount of memory the tracing system allocates in the server to contain trace records.

In cases where the memory buffer is full when a new trace record is created, the oldest entry from the memory buffer is deleted to make space.

Dump button

Use this button to write the contents of the trace memory buffer to the trace directory.

The tracing system resets the memory buffer after you dump it. The tracing system continues to record trace records in the memory buffer after you dump the buffer.

Directory to use for tracing and dumping memory buffer

Specifies the directory to which trace files are written. A subdirectory, trace data, is created in this directory, and the trace files are written to this location.

Note: This path cannot be shared between servers. The server creates a file with a .owner extension to help detect when two or more servers happen to be trying to use the same path for HPEL output.

Save runtime changes to configuration as well

Specifies that changes are made to both the dynamic state of the running server, and the server configuration, which takes effect on the next restart. If this check box is not selected, the server does not copy the settings into the server configuration.

HPEL text log configuration settings

Use this page to configure High Performance Extensible Logging (HPEL) settings for text log.

Note: You can only access this page when the server is configured to use HPEL log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and trace > *server_name* > Configure HPEL Text Log**.

Enable Text Log

Enables writing log and trace records into the text log file.

Specifies that in addition to writing log and trace records in binary format, the logging system writes them in a text format as well. You can configure the text log to be formatted in either of the formats that the basic mode SystemOut.log file uses.

Note: All content written to the text log is also written to either the log repository or trace repository. Enabling the text log degrades performance for applications that frequently create log or trace entries.

Directory path

Specifies the directory to which log files are written.

Text log file names have the following format: TextLog_<yy.mm.dd>_<hh.mm.ss>, where “TextLog_” is a fixed prefix, <yy.mm.dd> is a date (year, month, date) of the first record in the file, and <hh.mm.ss> is the time (hour, minute, second)

Note: This path cannot be shared between servers. The server creates a file with a .owner extension to help detect when two or more servers happen to be trying to use the same path for HPEL output.

Enable log record buffering

Specifies that the logging system avoids writing to disk each time a log record is created. The logging system creates a buffer that can hold a number of log records, and writes the buffered events when the buffer is full. The logging system also writes the buffered events after a few seconds have passed, even if the buffer is not full.

Selecting this setting significantly improves logging performance; however, if the server stops unexpectedly, the contents might not be written to the text log file.

Note: Enable log record buffering in almost all cases. Only disable log record buffering when your server is failing unexpectedly and cannot write buffered content to disk before stopping.

Start new log file daily at <time>

Enables the logging framework to close the log file and start a new file at the specified time of day. Closing the file makes it easy to copy the file to an archive.

Note: If you want to automatically archive your log files, set up your backup program to copy files after the time you configured for new logs to be started. Configure the backup to occur at least 10 minutes after the time configured for new logs to be started to ensure that the server has closed the previous file.

Begin cleanup of oldest records

Specifies the log cleanup settings to be used to automatically purge the oldest log records, or log records that no longer fit in the configured space, from the text log directory.

Select **When log size approaches maximum** to configure automatic log file cleanup to begin when the total size of the text log files approaches the configured maximum size.

Select **When oldest records reach age limit** to configure automatic log file cleanup to begin when log content is the age limit specified.

Select **When either age or size restriction is met** to configure automatic log file cleanup to begin when either of the previous conditions is met.

Regardless of the selection chosen, text log files are deleted from the text log directory in the order in which they were written.

Log record age limit:

Specifies the lifespan, in hours, that log records can remain in the text log directory before the log records can be automatically deleted by the server. When all records in a text log file have existed longer than the age limit specified, then that file is targeted for deletion by the server.

Maximum log size:

Specifies the maximum total size, in megabytes, that the server allows the text log files to reach. When the total size of the text log files approaches this size limit, the server deletes the oldest text log files from the text log directory to make space for new log records.

Out of space action

Specifies how the server reacts to an inability to add content to the text log directory.

Select **Stop server** to specify that the server stops when the server is unable to write to the text log directory.

Select **Purge old records** to specify that the server continues to run and that the file containing the oldest log records is immediately removed when the server is unable to write to the text log directory.

Select **Stop logging** to specify that the server continues to run, but that the server cannot continue to write to the log when the server is unable to write to the text log directory.

Text Output Format

Specifies the format to use in the text log file.

Select **Basic** to specify a shorter, one-line-per-record format.

Select **Advanced** to specify a longer format using full logger name and more details about each record.

Include trace records

Specifies whether trace records are included in the text log file, as well as log records.

Save runtime changes to configuration as well

Specifies that changes are made to both the dynamic state of the running server and the server configuration, which take effect on the next restart. If this check box is not selected, the server does not copy the settings into the server configuration.

Change log and trace mode settings

Use this page to choose the mode your system uses for logging and tracing.

Note: You can only access this page when the server is configured to use HPEL log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and trace > *server_name* > Change log and trace mode.**

Cancel



Use the Cancel button to keep the server in HPEL log and trace mode.

Switch back to basic mode



Use the **Switch to Basic Mode** button to change the log and trace mode for the server to basic.

Important: Switching the server to basic log and trace mode requires a server restart.

Log viewer settings

Use this page to view your High Performance Extensible Logging (HPEL) log, trace, System.out, and System.err content.

Note: You can only access this page when the server is configured to use HPEL log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and trace > *server_name* > View HPEL logs and trace.**

If HPEL is not enabled, you must enable it. Read about changing to HPEL logging and tracing from basic mode for more information.

You can only access this page when the server is configured to use HPEL log and trace mode. You can also view log and trace data for servers that are currently stopped as long as a node agent or administrative agent is running on the same machine and that agent is configured to use HPEL.

Log view table

Displays the log, trace, System.out, and System.err records.

The log view section displays the records. Use the First Page, Previous Page, Next Page, and Last Page buttons to move through the list of records, or specify filter criteria in the Content and Filtering Details section to limit the rows displayed. Records are always displayed in the order they were recorded by the server. By default the log view has 5 columns as listed in the following table..

Table 28. Log view table columns. The table columns provide information on the log records.

Column	Description
Time Stamp	The time when the event was recorded.
Thread ID	The identity of the thread that recorded the event in hexadecimal notation.
Logger	The logger that recorded the event.
Level	The type of event that was recorded.

Table 28. Log view table columns (continued). The table columns provide information on the log records.

Column	Description
Message	The message from the recorded event. If the message has a message ID the message ID is underlined. Click the message ID to get an explanation and recommended user action for the message.

To manipulate the log view, you can complete the following actions using available buttons.

Table 29. Log view table buttons. The table buttons help you configure the log settings.

Button	Resulting action
Refresh View	Clears the contents of the viewer and reinitializes the view using records from the server. Use this button to retrieve information about any additional rows created since the log viewer was started.
Show Only Selected Threads	Filters any records created by any thread other than the one selected in the selection area. Clicking on this button enables the Show All Threads button.
Show All Threads	Displays any records that were filtered when you clicked Show Only Selected Threads.. This button is only enabled when you have restricted the view using the Show Only Selected Treads button.
Select Columns	Enables you to select the columns in the viewer that you want to view.
Export	Exports logs to local workstation in any of basic, advanced, or binary (HPEL) formats
Copy to Clipboard	Copies the records that are highlighted in the selection area into the operating system clipboard.
Server Instance Information	Displays attributes for the selected server instance process. Use this table to find attributes and corresponding values for the server instance process environment. These properties are similar to the ones found in the header of basic mode logs.

Content and filtering details

Provides selection options to specify what content sources to include and what content to filter from the log view.

To change what content sources are shown in the table, or to filter what content is shown, expand the **Content and Filtering Details** section.

Server instance

Changes the server instance from which log records are retrieved.

A server instance represents a run of a server process. Each time the server is restarted, a new server instance is created. By default, the log view table shows log records generated in the most recent server instance. To select a different server start time, choose a server instance with the appropriate start timestamp. The timestamps shown represent the timestamp of the first record written to each server instance.

View contents

Controls what content sources are displayed in the log view.

System Out

Specifies that content logged to the System.out output stream is included in the log view.

System Error

Specifies that content logged to the System.err output stream is included in the log view.

Logs and trace

Specifies that log and trace records are included in the log view.

Log and trace entries can be further specified to include a minimum or maximum level. Minimum and maximum can be specified together, for example to display only a certain level of trace. If log and trace is not selected neither log nor trace records of any severity might be displayed.

Examples of log and trace filtering:

- Selecting logs and trace and clearing minimum level and maximum level fields results in the log view displaying records with any log or trace level.
- Selecting logs and trace and setting minimum level to WARNING results in the log view displaying log records with levels WARNING, FATAL, or SEVERE.
- Selecting logs and trace and setting maximum level to FINE results in the log view displaying trace records with levels FINE, FINER, or FINEST.
- Selecting logs and trace and setting minimum level to DETAIL and maximum level to AUDIT results in the log view displaying log records with levels DETAIL, CONFIG, INFO, or AUDIT.

Filtering

Controls which records are included in and excluded from the log view.

For all filters in this section, multiple entries might be specified using a colon (:) as a separator character. A limited set of regular expression characters can be used. Refer to console documentation for more details. If multiple filter settings are specified, the filter conditions must all be true for a record to be displayed in the log view.

Include loggers

Specifies the list of loggers whose records are included in the log view.

Exclude loggers

Specifies the list of loggers whose records are excluded from the log view.

Message contents

Specifies the message content that each record must contain to be included in the log view.

Event timing

Controls what records are displayed in the log view based on a start and end date and time.

From

Specifies the time of day, which the record creation time must be greater than or equal to for the record to be displayed in the log view.

Time must be specified as HH:mm:ss.SSS using the 24-hour clock. If the From value is not specified, a default value of 00:00:00.000 is used.

- HH represents the hour of the day. Valid values are from 0 to 23.
- mm represents the minute of the hour. Valid values are from 0 to 59.
- ss represents the seconds of the minute. Valid values are from 0 to 59.
- SSS represents the milliseconds of the second. Valid values are from 000 to 999.

On (first occurrence)

Specifies the date, which the record creation date must be greater than or equal to for the record to be displayed in the log view.

Until

Specifies the time of day, which the record creation time must be less than or equal to for the record to be displayed in the log view.

Time must be specified as HH:mm:ss.SSS using the 24-hour clock. If the Until value is not specified, a default value of 23:59:59.999 is used.

- HH represents the hour of the day. Valid values are from 0 to 23.
- mm represents the minute of the hour. Valid values are from 0 to 59.
- ss represents the seconds of the minute. Valid values are from 0 to 59.
- SSS represents the milliseconds of the second. Valid values are from 000 to 999.

On (second occurrence)

Specifies the date, which the record creation date must be less than or equal to for the record to be displayed in the log view.

LogViewer command-line tool

Use the LogViewer command to query the contents of the High Performance Extensible Logging (HPEL) log and trace repositories. You can also use the LogViewer command to view new log and trace repository entries as the server writes content to them.

LogViewer

The High Performance Extensible Logging (HPEL) facility writes to the log and trace repositories in a binary format. You can view, query and filter the repository using the LogViewer command. The LogViewer command provides options for quickly converting HPEL logs into a text file in various formats, including basic, advanced, and Common Base Event format. The command also provides options to make getting the data you need from the logs easier; for example, allowing you to filter what log records you want by level, logger name, or date and time.

Use the following command to view the full contents of your log and trace repositories:

Optional parameters

-repositoryDir *directory_name*

Specifies the path to the repository directory. In the case where you want to query both the log and trace data together, provide the path to the parent directory, which contains both the log data and tracedata directories. If you use the default repository location, `profile_root/logs/application_server/`, and run this tool from the profile bin directory, then this argument is optional. The tool checks the default location if one is not provided. If multiple application servers exist in this profile with HPEL repositories, you are prompted to select which server log and trace repository you want to view.

-outLog *file_name*

Specifies the file name you want the text output written to. If you do not provide this information, the text output is displayed on the console.

-format **basic** | **advanced** | **cbe-1.0.1**

Specifies the output format. Supported formats include basic, advanced, and the CBE-1.0.1 format. If you do not provide this information, the output is in basic format.

- monitor** *[integer]*
Specifies that you want the logViewer to continuously monitor the repository and output new log record entries as they are created. You can provide an optional integer argument after this parameter to specify how often you want the LogViewer tool to query the repository for new records. By default the logViewer queries the repository for new records every 5 seconds. When used with other filtering options, only those new records that match the filter criteria are displayed.
- help**
Use this parameter to have the LogViewer tool list the full set of options that are available.
- startDate** *date_time*
You can filter the results that are displayed from the repository by date and time. Use the startDate parameter to filter out log entries that occurred before the date or date time provided as an argument. Provide either a date or date and time, entered in the MM/dd/yy format or the MM/dd/yy H:m:s:S format.
- stopDate** *date_time*
Use this parameter to filter out log entries that occurred after the specified date or date time. Provide the argument in the same format as the -startDate option.
- level** *level_name*
Specifies that you want the tool to only display those log events which match the level name you provide as an argument. Valid values for the level name are FINEST, FINER, FINE, DETAIL, CONFIG, INFO, AUDIT, WARNING, SEVERE, FATAL.
- minLevel** *level_name*
Specifies that you want the tool to only display records which are at or above the specified level. Valid values for the level name are FINEST, FINER, FINE, DETAIL, CONFIG, INFO, AUDIT, WARNING, SEVERE, FATAL.
- maxLevel** *level_name*
Specifies that you want the tool to only display records that are at or below the specified level. Valid values for the level name are FINEST, FINER, FINE, DETAIL, CONFIG, INFO, AUDIT, WARNING, SEVERE, FATAL.
- includeLoggers** *logger_name*
When this option is used, only log events from the specified loggers are included in the LogViewer output. Separate multiple entries with a comma. The * symbol can be used as a wild card to include all loggers below a parent logger. When used in combination with the -excludedLoggers option, the more specific match determines if the log event is included or excluded.
- excludeLoggers** *logger_name*
Use this option to exclude log events from the specified loggers in the LogViewer output. Separate multiple entries with a comma. The * symbol can be used as a wildcard to include all loggers below a parent logger. When used in combination with the -includeLoggers option, the more specific match determines if the log event is included or excluded.
- thread** *thread_id*
Use this option to restrict LogViewer output to only those log events from a specific thread. Any log messages that were not created by the thread ID provided as an argument to this option are not displayed. Specify the thread ID in hex format.
- extractToNewRepository** *directory_name*
This option redirects log and trace records from a binary repository to a new binary repository at the location that you specify. You can use this option with other filtering options to get a subset of log and trace records into the new repository. This option uses the directory path where the new repository must be written as an argument. Therefore, the directory must be empty. If the directory does not exist, the directory is created. However, errors that occur during the directory creation might create extraneous directories.

-listInstances

Use this option to list the IDs of available server process instances that are available to use with the `-instance` option. After running LogViewer with the `-listInstances` option, you can then use the `-instance` option to invoke LogViewer with one of the server process instance IDs as an argument. Since this option does not process any log or trace records, all other options are ignored when you specify this option.

-instance *instance_id*

Use this option to retrieve the log and trace data for a given server process instance by providing the server instance ID. Run LogViewer, along with the `-listInstances` option, before you use this option to obtain a valid instance ID. This option is required when viewing logs and trace from an environment that contains subprocesses, such as the z/OS operating system.

If this option is combined with `-latestInstance`, `-instance` is ignored.

-latestInstance

Use this option to retrieve the log and trace data from the most recent server instance. If this option is used with the `-instance` option, the `-instance` option is ignored.

-message *match_string*

Use this option to retrieve only log or trace data with a message field that matches the requested text.

-includeExtensions *name*[=*value*][,*name*[=*value*]]*

Use this option to retrieve the log and trace data with an extension name that matches the requested name, and an extension value that matches the requested value. You can also use this option to retrieve the log and trace data with an extension name that matches the requested name, and an extension value that matches any value, if you omit the `=value` part of the option.

Any extension name shown in the advanced format can be used. Note that 'source', 'class', and 'method' are not stored in the log/trace repositories as extensions, and so cannot be filtered on with this option.

Separate multiple `name=value` arguments with a comma. Specify '==' (two equals signs) in place of '=' (one equals sign) in cases where the name or value must contain an equal sign. Specify ',' (two commas) in place of ',' (one comma) in cases where the name or value must contain a comma.

-encoding *character_set*

Specifies the character set that the LogViewer command will use for text output.

Filtering considerations

Be aware of LogViewer filtering optimizations. The LogViewer tool is able to filter log and trace data most efficiently when used with the following filter options:

- `startDate`
- `stopDate`
- `thread`
- `level`
- `minLevel`
- `maxLevel`

Example usage

See the following examples of LogViewer commands on UNIX-based systems. The examples show how to run LogViewer from the profile bin directory where the `repositoryDir` parameter is not required.

- Write all records in the default repository between July 19th, 2009 and August 2nd, 2009 to a file called `/tmp/promo.logs`.

```
logViewer.sh -outLog /tmp/promo.logs -startDate 07/19/2009 -stopDate 08/02/2009
```

- Display new records whose specified level is WARNING or higher using the advanced format as the server writes them to the log repository.

```
logViewer.sh -monitor -minLevel WARNING -format advanced
```

- Write only those log messages that were written to the error stream of a specific repository to a file called logged_errors.txt.

```
logViewer.sh -repositoryDir /apps/server1/logs -includeLoggers SystemErr -outLog logged_errors.txt
```

- View events from the default repository that occurred before September 14th, 2009 4:28 PM eastern daylight time.

```
logViewer.sh -stopDate "09/14/2009 16:28:00:000 EDT"
```

- Write events from the default repository that contain a 'thread' extension with value 'WebContainer : 6'

```
logViewer.sh -includeExtensions thread="WebContainer : 6" -format advanced
```

- Write events from the default repository that were a part of the request with requestID a856cb2c-79ed-4d62-a3cf-a9908b2db07b.

```
logViewer.sh -includeExtensions requestID=a856cb2c-79ed-4d62-a3cf-a9908b2db07b
```

- Write events from the default repository that were created on a thread servicing the PlantsByWebSphere application.

```
logViewer.sh -includeExtensions appName=PlantsByWebSphere
```

Monitoring application logging using JMX notifications

Java developers can create programs to monitor application server logs using JMX notifications.

About this task

The most common log message listeners are written in Java, and connect to the deployment manager or an application server using SOAP. Use this topic to build a Java client that listens for log events.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log, SystemErr.log, trace.log, and activity.log files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Note: Be careful when adding listeners to servers with high logging volume as JMX notifications can slow down your server.

Procedure

1. Import the necessary packages. You will typically need the following import statements at the beginning of your Java program:

```
import javax.management.Notification;
import javax.management.NotificationListener;
import javax.management.ObjectName;
import javax.management.InstanceNotFoundException;
import javax.management.MalformedObjectNameException;
```

```
import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.exception.ConnectorException;
```

Additionally, to handle the messages, and the types returned from the calls in subsequent steps you will need the following import statements.

```
import java.util.Iterator;
import java.util.Properties;
import java.util.Set;
import com.ibm.websphere.ras.RasMessage;
```

2. Create a Java class that implements the NotificationListener interface.
3. Implement the handleNotification method. The following example is a sample that writes the message text to the Java console:

```
public void handleNotification(Notification notification, Object handback) {
    RasMessage rasMessage = (RasMessage)notification.getUserData();
    System.out.println("Localized message: " + rasMessage.getLocalizedMessage(null));
}
```

4. Connect to the SOAP port of the server whose JMX MBeans you want to monitor. The following code creates a SOAP-connected AdminClient object with a specified host and a specified port:

```
AdminClient adminClient = null;
String hostName = "someHostName";
String soapPort = "8880";

Properties connectProps = new Properties();
connectProps.setProperty(AdminClient.CONNECTOR_TYPE, "SOAP");
connectProps.setProperty(AdminClient.CONNECTOR_HOST, hostName);
connectProps.setProperty(AdminClient.CONNECTOR_PORT, soapPort);

try {
    adminClient = AdminClientFactory.createAdminClient(connectProps);
} catch (ConnectorException e) {
    // error handling code
}
```

5. Retrieve the MBean object name for the RasLoggingService MBean. The following code retrieves the RasLoggingService MBean object name:

```
String queryString = "WebSphere:cell="+cellName+",node="+nodeName+",process="+serverName+",
type=RasLoggingService,*" ;
Set<ObjectName> objectMBeans = null;
try {
    ObjectName queryName = new ObjectName(queryString);
    objectMBeans = (Set<ObjectName>)adminClient.queryNames(queryName, null);
    if (objectMBeans.size() > 1) {
        // error handling code to deal with the case where we get more than one name returned.
    }
} catch (MalformedObjectNameException e) {
    // error handling code
} catch (ConnectorException e) {
    // error handling code
}

if (objectMBeans.isEmpty()) {
    // error handling code to deal with the case where the MBean is not found
}

Iterator<ObjectName> objectNames = objectMBeans.iterator();
ObjectName objectName = objectNames.next();
```

6. Add the notification listener. This sample code adds a notification listener, waits for 60 seconds while it processes notifications, then removes the notification listener. A listener can stay connected as long as needed.

```
try {
    adminClient.addNotificationListener(objectName, this, null, null);
    Thread.sleep(60 * 1000);
    adminClient.removeNotificationListener(objectName, this);
} catch (InstanceNotFoundException e) {
    // error handling code
} catch (ConnectorException e) {
    // error handling code
} catch (Exception e) {
    // error handling code
}
```

7. Add the necessary jar to your classpath. Add the admin client jar file to your classpath to be able to compile and run your code. The admin client jar file is in the <install_root>/runtimes directory.

Results

You have created a Java program that can listen to, and take actions as a result of log event notifications from an application server.

Chapter 7. Using Cross-Component Trace to troubleshoot applications

You can use Cross-Component Trace (XCT) to help diagnose problems in WebSphere Application Server.

About this task

Administrators using WebSphere Application Server need to use log and trace files to determine whether their applications and the server are running correctly.

Depending on the nature of your applications, multiple threads within an application server may be used to handle requests, such as HTTP requests or JMS requests. Some requests may be handled by more than one application server, such as when one application server makes a request to another application server for a web services request.

You can use XCT to augment your log and trace files with correlation information. This correlation information clarifies which threads and which application server processes participated in the handling of each request.

Procedure

1. Enable XCT if you have not done so already.
 - a. Log on to the administrative console.
 - b. If using an admin agent topology, select a node that you want to manage and navigate to it.
 - c. From the navigation section in the console, choose **Troubleshooting > Logs and trace**.
 - d. Select the server that you want to enable XCT for.
 - e. Choose **Change log detail levels**.
 - f. Select the **Configuration** tab if you want to make static configuration changes, or the **Runtime** tab if you want to make changes to the runtime state of the server.
 - g. Check **Enable log and trace correlation** checkbox.
 - h. Select **Include request IDs in log and trace records, Include request IDs in log and trace records and correlation log records**, or **Include request IDs in log and trace records, create correlation log records, and capture data snapshots** as appropriate.
 - i. Click **OK**.
 - j. If you made your changes on the Configuration tab, save them and restart the application server.


Note: Enable XCT to include request IDs in log and trace files when you want to see which log and trace entries, in all threads and application server processes, are related to the same request. Request IDs are only recorded when using HPEL log and trace mode and can be seen or used for filtering using the logViewer command.


Note: Enable XCT to create correlation log records when you want to log how requests branch between threads and processes, and see extra information about each request. Enabling XCT to create correlation log records might have a significant performance impact on your system, so is best suited to test and development environments.

Note: Enable XCT to capture data snapshots when you want to store entire request and response bodies to the file system. Enabling XCT to capture data snapshots might have a significant performance impact on your system, so is best suited to test and development environments. XCT captures data snapshots for message requests and responses handled by the SIBus.

Note: Data snapshots are captured and written to the `$SERVER_LOG_ROOT/snapdata` directory. The application server does not automatically clean up files from this directory. You will need to


delete the files from this directory periodically when data snapshot capturing is enabled. Data snapshots store entire request and response contents and may include sensitive information. This option might not be appropriate for use in production environments.

2. Use XCT request ID information to track requests.
 - a. Ensure you are using High Performance Extensible Logging (HPEL) log and trace mode, since basic mode log and trace does not store request IDs. Read the Changing from basic mode to HPEL logging and tracing topic for more information.
 - b. Enable XCT if you have not done so already.
 - c. Filter your logs to look for important information, such as errors or warnings using the HPEL LogViewer command-line tool. Output your logs using advanced format so that you can see the request ID information in the logs. For example, use the command, as follows: 

```
logViewer.sh -minLevel WARNING -format advanced
```
 - d. When you have found log entries that are of interest to you, note the request ID associated with those entries.
 - e. Filter your logs by request ID using the HPEL LogViewer command-line tool and using the request IDs you noted in the previous step, as follows: 

```
logViewer.sh -includeExtensions requestID=<requestID>
```

Read the LogViewer command-line tool topic for more information.

3. Use XCT correlation log records to determine the hierarchy of call chains.
 - a. Enable XCT if you have not done so already. Select **Include request IDs in log and trace records and correlation log records** to ensure that XCT creates correlation log records.
 - b. If you are using HPEL, convert your logs to text using the LogViewer command-line tool. For example, use the command, as follows: 

```
logViewer.sh -outLog myLog.txt
```
- Read the LogViewer command-line tool topic for more information.
- c. Use the XCT Log Viewer tool available with the IBM Support Assistant to load log and trace files from all application servers involved in handling your requests and see the hierarchy of request call chains. Read the Using IBM Support Assistant topic for more information.

Results

The server is now configured to use XCT.

Cross Component Trace (XCT)

Cross Component Trace (XCT) annotates the logs so that log entries that are related to a request that is serviced by more than one thread, process, or even server are identified as belonging to the same unit of work.

Overview

Note: XCT helps identify the root cause of problems across components, which provides the following benefits:

- Enables administrators and support teams to follow the flow of a request from end-to-end as it traverses thread or process boundaries, or travels between stack products and WebSphere Application Server.
- Helps to resolve questions about which component is responsible for a request that fails.

Administration of XCT

XCT is a capability built into the WebSphere Application Server log and trace framework. When enabled, XCT annotates the logs so that log entries that are related to a request that is serviced by more than one thread, process, or even server are identified as belonging to the same unit of work. Applications built using distributed architectures, such as Service Oriented Architecture, will benefit from XCT, since XCT helps facilitate problem determination across multiple services on different systems.

XCT different modes:

- Fully disabled.
- With XCT request IDs added to existing log and trace records.
- With XCT request IDs added to existing log and trace records and XCT log records added to log files.
- With XCT request IDs added to existing log and trace records, XCT log records added to log files, and data snapshots captured.

XCT request IDs:

- XCT request IDs are identifiers added to log and trace records produced by the server when the server is configured to use High Performance Extensible Logging (HPEL).
- XCT adds the same request ID to every log or trace record as long as the log or trace record is a part of the same request, regardless of which thread or Java virtual machine (JVM) produces the log or trace entry.
- When XCT is used with the HPEL log and trace infrastructure, you can view request IDs with the logViewer tool when logs are output in advanced format.

The following is an example of a log record with an XCT request ID in the log file (shown rendered in advanced format):

```
[3/18/11 14:50:17:391 EDT] 00000018 W UOW= source=com.ibm.somelogger.QuickLogTest org= prod= component=
thread=[WebContainer : 1] requestID=BJrcVPo+Yk4-AAAAAAA8zAA hello world
```

Note that the request ID is shown previously on a separate line, but in the log files it is actually on the same line as the rest of the log record header.

XCT log records:

- XCT log records are typically added to the logs to:
 - demarcate the beginning and ending of work for a particular request on a particular thread.
 - demarcate when work is about to be transferred to another thread or process, or to indicate when work returned from another thread or process.
 - demarcate when work moves from major component to major component, even if work continues on the same thread; for example to show transfer of control from application server code to application code.

The following is an example of an XCT log record in the log file:

```
[3/18/11 14:50:17:391 EDT] 00000031 XCT I BEGIN BJrcVPo+Yk4-AAAAAAA8zAA 000000000000-ccccccccc2
HTTPCF(OutboundRequest /index.html RemoteAddress(127.0.0.1) RequestContext(36001645))
```

- XCT log records are composed of:
 - XCT type (BEGIN / END)
 - XCT parent correlator ID (for example, 000000000000-ccccccccc2)
 - XCT current correlator ID (for example, BJrcVPo+Yk4-AAAAAAA8zAA)
 - XCT annotations (for example, HTTPCF(OutboundRequest /index.html RemoteAddress(127.0.0.1) RequestContext(36001645))

XCT tools:

- The HPEL logViewer tool is able to filter log and trace records by request ID.

- Tools, such as the XCT Log Viewer, can also take advantage of XCT log records or XCT request IDs, or both, when rendering log and trace content. The XCT Log Viewer is available as a tool add-on for the IBM Support Assistant.

XCT configuration:

You can configure XCT using wsadmin scripting.

Configuring XCT with wsadmin scripting

You can configure Cross-Component Trace (XCT) using wsadmin scripting. Use the examples in this topic as a guide to build your own wsadmin scripts.

About this task

XCT annotates the logs so that log entries that are related to a request that is serviced by more than one thread, process, or even server are identified as belonging to the same unit of work. You can configure XCT using the administrative console, or using wsadmin scripting. The examples in this topic show how to configure XCT using wsadmin. If you complete this task using the deployment manager, then you might need to synchronize the node agent on the target node and restart the server before configuration changes take effect.

Table 30. Variable Names. The table applies to all examples in this topic. All examples use the Jython scripting language.

Variable	Description
<i>myCell</i>	The name of the cell
<i>myNode</i>	The host name of the node
<i>myServer</i>	The name of the server

Procedure

- Use the AdminConfig object to configure XCT.

Changes you make using the AdminConfig object take effect the next time you start the server.

1. Enable XCT for the server.

By default, XCT is disabled for the server. The following example shows how to enable XCT for the server when your server is using High Performance Extensible Logging (HPEL) log and trace mode. Specify RASLoggingService instead of HighPerformanceExtensibleLogging when your server is using basic log and trace mode. Specify false instead of true to disable XCT.

```
# get rid of existing property if already present
configId = AdminConfig.getid("/Cell:myCell/Node:myNode
/Server:myServer/HighPerformanceExtensibleLogging:
/Property:com.ibm.websphere.logging.enableCorrelation")if (len(configId) > 0):
    AdminConfig.remove(configId)
# add new property
LoggingService = AdminConfig.getid("/Cell:myCell/Node:myNode
/Server:myServer/HighPerformanceExtensibleLogging:")
AdminConfig.create("Property", LoggingService, [{"name", "com.ibm.websphere.logging.enableCorrelation"}, {"value", "true"}])

AdminConfig.save()
```

2. Change the log setting for XCT .

By default, the XCT log setting includes request IDs in log and trace files. The following example shows how to make XCT both include request IDs in log and trace files and log XCT log records. Specify REQUEST ID instead of LOG to make XCT only include request IDs in log and trace files. Specify DATA_SNAPSHOT instead of LOG to make XCT include request IDs in log and trace files, log XCT log records, and capture data snapshots. Specify RASLoggingService instead of HighPerformanceExtensibleLogging when your server is using basic log and trace mode.

```

# get rid of existing property if already present
configId = AdminConfig.getid("/Cell:myCell/Node:myNode
/Server:myServer/HighPerformanceExtensibleLogging:/Property:com.ibm.websphere
.logging.correlationLevel") if (len(configId) > 0): AdminConfig.remove(configId)
# add new property
LoggingService = AdminConfig.getid("/Cell:myCell/Node:myNode
Server:myServer/HighPerformanceExtensibleLogging:/")
AdminConfig.create("Property", LoggingService, [{"name", "com.ibm.websphere.logging.correlationLevel"}, {"value", "LOG"}])

AdminConfig.save()

```

Note: Enable XCT to include request IDs in log and trace files when you want to see which log and trace entries, in all threads and application server processes, are related to the same request. Request IDs are only recorded when using HPEL log and trace mode and can be seen or used for filtering using the logViewer command.

Note: Enable XCT to create correlation log records when you want to log how requests branch between threads and processes, and see extra information about each request. Enabling XCT to create correlation log records might have a significant performance impact on your system, so is best suited to test and development environments.

Note: Enable XCT to capture data snapshots when you want to store entire request and response bodies to the file system. Enabling XCT to capture data snapshots might have a significant performance impact on your system, so is best suited to test and development environments. XCT captures data snapshots for message requests and responses handled by the SIBus.

Note: Data snapshots are captured and written to the `$SERVER_LOG_ROOT/snapdata` directory. The application server does not automatically clean up files from this directory. You will need to delete the files from this directory periodically when data snapshot capturing is enabled. Data snapshots store entire request and response contents and may include sensitive information. This option might not be appropriate for use in production environments.

- Use the AdminControl object to configure XCT. Changes you make using the AdminControl object take effect immediately.

1. Enable XCT for the server.

By default, XCT is disabled for the server. The following example shows how to enable XCT for the server when your server is using HPEL log and trace mode. Ensure the server is running and specify RasLoggingService instead of HPELControlService when your server is using basic log and trace mode. Specify false instead of true to disable XCT.

```

LoggingMBean = AdminControl.queryNames('cell=myCell,node=myNode,
type=HPELControlService,process=myServer,*')
AdminControl.setAttribute(LoggingMBean, "correlationEnabled", "true")

```

2. Change the log setting for XCT .

By default, the XCT log setting includes request IDs in log and trace files. The following example shows how to make XCT both include request IDs in log and trace files and log XCT log records. Specify REQUEST ID instead of LOG to make XCT only include request IDs in log and trace files. Specify DATA_SNAPSHOT instead of LOG to make XCT include request IDs in log and trace files, log XCT log records, and capture data snapshots. Ensure the server is running and specify RasLoggingService instead of HPELControlService when your server is using basic log and trace mode.

```

LoggingMBean = AdminControl.queryNames('cell=myCell,node=myNode,type=HPELControlService,process=myServer,*')
AdminControl.setAttribute(LoggingMBean, "xctLevel", "LOG")

```

Note: Enable XCT to include request IDs in log and trace files when you want to see which log and trace entries, in all threads and application server processes, are related to the same request. Request IDs are only recorded when using HPEL log and trace mode and can be seen or used for filtering using the logViewer command.

Note: Enable XCT to create correlation log records when you want to log how requests branch between threads and processes, and see extra information about each request. Enabling XCT to create correlation log records might have a significant performance impact on your system, so is best suited to test and development environments.

Note: Enable XCT to capture data snapshots when you want to store entire request and response bodies to the file system. Enabling XCT to capture data snapshots might have a significant performance impact on your system, so is best suited to test and development environments. XCT captures data snapshots for message requests and responses handled by the SIBus.

Note: Data snapshots are captured and written to the `$SERVER_LOG_ROOT/snapdata` directory. The application server does not automatically clean up files from this directory. You will need to delete the files from this directory periodically when data snapshot capturing is enabled. Data snapshots store entire request and response contents and may include sensitive information. This option might not be appropriate for use in production environments.

Results

XCT is now configured. If you made changes with the AdminConfig command, restart the server to make the changes take effect.

Chapter 8. Using sensitive log and trace guard

You can protect information with the sensitive log and trace guard. The sensitive log and trace guard prevents loggers from writing sensitive information in your log and trace files.

Sensitive log and trace guard

The sensitive log and trace guard is a feature that helps administrators prevent sensitive information from being exposed in log and trace files.

The sensitive log and trace guard uses an internal list of allowable levels for sensitive loggers which specifies the lowest level at which listed loggers can generate log or trace data without containing potentially sensitive data. You can also add your own loggers to the list that the sensitive log and trace guard will block.

An example is as follows: If a servlet writes URL request parameters verbatim to logger `com.xyz.SomeLogger` at level `Level.FINE`, and these request parameters could contain information such as credit card numbers or passwords, then you should add an entry to the sensitive logger list to allow only levels higher than `Level.FINE` to be logged - `com.xyz.SomeLogger=CONFIG`.

When the server initializes the log and trace system, or when you attempt to change the log detail level for a server, the list of allowable levels for sensitive loggers is compared to the stated log detail level. Any attempt to enable logging or tracing that is in conflict with entries in the list is overridden. In cases where the same loggers are specified multiple times in the list the most restrictive entry is used. For example, if the list of allowable levels for sensitive loggers contains `a.b.*=INFO` and `a.b.*=FINE`, then the `a.b.*` loggers are restricted to only being able to log at levels `INFO` and higher. The following table provides examples of how lists of allowable levels for sensitive loggers modify specified log detail level settings to determine effective log detail level settings:

Table 31. Restriction list. Examples of how sensitive log and trace guard settings affect log detail level settings

Log detail level setting	List of allowable levels for sensitive loggers	Resulting effective log detail level setting
<code>a.b.*=SEVERE</code>	<code>a.b.*=FINE</code>	<code>*=INFO:a.b.*=SEVERE</code>
<code>a.b.*=SEVERE</code>	<code>a.b.*=SEVERE</code>	<code>*=INFO:a.b.*=SEVERE</code>
<code>a.b.*=FINE</code>	<code>a.b.*=FATAL</code>	<code>*=INFO:a.b.*=FATAL</code>
<code>a.*=SEVERE</code>	<code>a.b.*=FINE</code>	<code>*=INFO:a.*=SEVERE</code>
<code>a.*=SEVERE</code>	<code>a.b.*=SEVERE</code>	<code>*=INFO:a.*=SEVERE</code>
<code>a.*=FINE</code>	<code>a.b.*=FATAL</code>	<code>*=INFO:a.*=FINE:a.b.*=FATAL</code>
<code>a.b.*=SEVERE</code>	<code>a.*=FINE</code>	<code>*=INFO:a.b.*=SEVERE</code>
<code>a.b.*=SEVERE</code>	<code>a.*=SEVERE</code>	<code>*=INFO:a.b.*=SEVERE</code>
<code>a.b.*=FINE</code>	<code>a.*=FATAL</code>	<code>*=INFO:a.b.*=FATAL</code>
<code>a.b.*=FINE</code>	<code>*=SEVERE</code>	<code>*=SEVERE:a.b.*=SEVERE</code>
<code>a.b.*=FINE</code>	<code>*=FINE</code>	<code>*=INFO:a.b.*=FINE</code>
<code>a.b.*=FINE</code>	<code>*=FINEST</code>	<code>*=INFO:a.b.*=FINE</code>
<code>a.b.*=FINE</code>	<code>x.y.z.*=SEVERE</code>	<code>*=INFO:a.b.*=FINE:x.y.z.*=SEVERE</code>
<code>a.b.*=FINE</code>	<code>x.y.z.*=FINE</code>	<code>*=INFO:a.b.*=FINE</code>
<code>a.b.*=FINE</code>	<code>x.y.z.*=FINEST</code>	<code>*=INFO:a.b.*=FINE</code>
<code>a.b.*=FINE</code>	<code>*=WARNING:x.y.z.*=SEVERE</code>	<code>*=WARNING:a.b.*=WARNING:x.y.z.*=SEVERE</code>
<code>a.b.*=FINE</code>	<code>*=WARNING:*=SEVERE:x.y.z.*=SEVERE</code>	<code>*=SEVERE:a.b.*=SEVERE</code>

By using this log and trace guard, you can prevent loggers from logging at levels which might expose sensitive information.

The product is preconfigured with a known list of loggers to restrict, however you might find that further restrictions are required. .

Enabling and disabling sensitive log and trace guard

You can either enable or disable the sensitive log and trace guard to help control whether loggers write sensitive information in your log and trace files.

About this task

Administrators using WebSphere Application Server can prevent sensitive information, such as data provided from users in HTTP requests, from being written in log and trace files. In some cases, when having access to private data can help with debugging, you might want to disable sensitive log and trace guard. For example, you might see that a credit card number that was entered in a web form did not have the required number of digits.

Sensitive log and trace guard works by preventing administrators from enabling certain loggers to levels at which they are known to log or trace sensitive information.

Use the administrative console to enable or disable the sensitive log and trace guard.

Procedure

1. Log on to the administrative console.
2. If you are using an administrative agent topology, then select a node that you want to manage, and navigate to it.
3. From the navigation section in the console, choose **Troubleshooting > Logs and trace**.
4. Select the server that you want to enable or disable with sensitive log and trace guard.
5. Click **Change log detail levels**.
6. Select the **Disable logging and tracing of potentially sensitive data** check box to enable sensitive log and trace guard. To disable sensitive log and trace guard, clear the **Disable logging and tracing of potentially sensitive data** check box.
7. Click **OK**
8. Save the changes.

Results

After you enable sensitive log and trace guard, the server is now configured to prevent known sensitive loggers from writing sensitive content to the log and trace files. After you disable sensitive log and trace guard, the server is now configured to allow known sensitive loggers to write sensitive content to the log and trace files. If you completed these steps using the deployment manager, you might need to synchronize the node agent on the target node before restarting the server.

Maintaining sensitive log and trace guard lists

The sensitive log and trace guard relies on lists which declare which loggers can potentially log or trace sensitive information, and the levels at which the sensitive information would be logged. You can extend the default list of loggers and their corresponding levels in cases where you find sensitive information in your log or trace that you want to block from being logged or traced in the future.

Before you begin

Read about log level settings for information about enabling the Sensitive Log and Trace Guard.

About this task

The application server has a private default list of sensitive loggers and their corresponding levels which it will block whenever the sensitive log and trace guard feature is enabled. The application server also provides a sensitive log and trace guard property file, and a sensitive log and trace guard API that you can use to declare new logger restrictions if you discover other loggers which log or trace sensitive information.

Note: If you attempt to add loggers to the sensitive log and trace guard list that have already been declared, the sensitive log and trace guard will use the more restrictive logger setting of the already declared and newly specified levels. For example, if the server is already configured to only allow logger `com.xyz.SomeLogger` to log at level `FINE`, and you attempt to declare that the same logger should only be allowed to log at level `FINEST`, the server will ignore the update, but if you attempt to declare that the same logger should only be allowed to log at level `INFO`, then the server will reconfigure the sensitive log and trace guard to use level `INFO` for that logger.

Procedure

- You can use a properties file to declare new logger restrictions. This file is in the cell-scoped configuration for each profile. The name is:

```
<profileHome>/config/cells/<cellname>/ras.rawtracelist.properties
```

This file contains documentation and syntax samples, but contains no actual entries. If you edit this file on the deployment manager the file is automatically synchronized with all nodes in the cell. If you edit this file on a specific node, it will be replaced the next time the file is synchronized with the deployment manager. Thus, it is best to maintain the list at the deployment manager.

- You can use the `com.ibm.websphere.logging.RawTraceList` API to declare new logger restrictions. This API allows you to add individual entries or an array of entries (using the `PatternLevel` object in the same package). It also allows passing in an input stream in the same format as the properties file.



Chapter 9. Diagnosing problems (using diagnosis tools)

Various diagnosis tools are provided to help you determine the source and impact of problems occurring in your application serving environment.

About this task

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 non-optimized performance, problem determination focuses on finding solutions to functional problems.

Procedure

1. If deploying or running an application results in exceptions such as `ClassNotFoundException`, use the Class Loader Viewer to diagnose problems with class loaders.
2. If you already have an error message and want to quickly look up its explanation and recommended response, look up the message by expanding the Messages section of the Information Center under Reference > Messages.
3. If you are using the basic or traditional log and trace infrastructure, and need help finding error and warning messages, interpreting messages, and configuring log files, see Chapter 10, “Using basic or traditional message logs to troubleshoot applications,” on page 139.
4. If you are using the High Performance Extensible Logging (HPEL) log and trace infrastructure, and need help finding error and warning messages, interpreting messages, and configuring log files, see Chapter 6, “Using High Performance Extensible Logging to troubleshoot applications,” on page 97.
5. Difficult problems can require the use of tracing, which exposes the low-level flow of control and interactions between components. For help in understanding and using traces, see *Working with trace*.
6.  For help in using settings or tools to help you diagnose the problem, see *Working with troubleshooting tools*. Some of these tools are bundled with the product, and others are freely downloadable.
7. To learn how to work with Diagnostic Providers, see *Working with Diagnostic Providers*.
8. To find out how to look up documented problems, common mistakes, WebSphere Application Server prerequisites, and other problem-determination information on the WebSphere Application Server public website, or to obtain technical support from IBM, see *Obtaining help from IBM*.
9.  The *Troubleshoot IBM Developer Kit for Java* 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.
10. For current information available from IBM Support on known problems and their resolution, see the WebSphere Application Server Product support page. For last minute updates, limitations, and known problems, refer to the Release notes section.
11. Before opening a PMR, there is MustGather information that you need to collect to send to IBM Support. This information assists IBM Support in identifying and resolving your problem. You can use the IBM Support Assistant Data Collector to automatically collect the MustGather information from your system.
12. To automatically gather Must gather type information from your system, see *IBM Support Assistant Data Collector*.


Chapter 10. Using basic or traditional message logs to troubleshoot applications

WebSphere Application Server can write system messages to several general purpose logs, including JVM, process, and IBM service logs, which can be examined for problem determination.

Before you begin

The JVM logs are created by redirecting the `System.out` and `System.err` streams of the JVM to independent log files. WebSphere Application Server writes formatted messages to the `System.out` stream. In addition, applications and other code can write to these streams using the `print()` and `println()` methods defined by the streams. Some Developer Kit built-ins such as the `printStackTrace()` method on the `Throwable` class can also write to these streams. Typically, the `System.out` log is used to monitor the health of the running application server. The `System.out` log and `System.err` log can be used for problem determination. The `System.err` log contains exception stack trace information that is useful when performing problem analysis.

Because each application server represents a JVM, there is one set of JVM logs for each application server and all of its applications located by default in the following directory:

-  `profile_root/logs/server_name`

The process logs are created by redirecting the `STDOUT` and `STDERR` streams of the process to independent log files. Native code, including the Java virtual machine (JVM) itself, writes to these files. As a general rule, WebSphere Application Server does not write to these files. However, these logs can contain information relating to problems in native code or diagnostic information written by the JVM.

As with JVM logs, there is a set of process logs for each application server, since each JVM is an operating system process.

Note: The IBM service log contains both the WebSphere Application Server messages that are written to the `System.out` stream and some special messages that contain extended service information that is normally not of interest, but can be important when analyzing problems. There is one service log for all WebSphere Application Server JVMs on a node, including all application servers. The IBM Service log is maintained in a binary format and requires a special tool to view. This viewer, the Log and Trace Analyzer, provides additional diagnostic capabilities. In addition, the binary format provides capabilities that are utilized by IBM support organizations.


In addition to these general purpose logs, WebSphere Application Server contains other specialized logs that are specific to a particular component or activity. For example, the HTTP server plug-in maintains a special log. Normally, these logs are not of interest, but you might be instructed to examine one or more of these logs while performing specific problem determination procedures. For details on how and when to view the plug-in log, see the [Accessing a web resource through the application server and bypassing the HTTP server subsection of the A web resource does not display topic](#).

About this task

Sometimes server and application problems can be diagnosed by examining log output from the WebSphere Application Server.

Procedure

Determine which type of logs you would like to implement:

- JVM logs
-  Process logs

- IBM service logs

Example

How to direct SYSPRINT and SYSOUT output to an HFS file.

If you are familiar with UNIX or Windows environments, you might be reluctant to use the facilities of SDSF (or IOF) to view the SYSPRINT and SYSOUT output from servants. If you would rather use a familiar editor (such as vi) in a Telnet session to view your output, it is possible to redirect the SYSPRINT and SYSOUT outputs to files in an HFS.

The JCL example below shows how to modify the SYSPRINT DD card in your startup procedure to redirect the output to an HFS file. The old SYSPRINT DD card has been commented out by preceding it with /*, and a new SYSPRINT DD card points to a file in the "/myDir/myServer" directory, in this case named was.log.d&LYMMDD..t&LHHMSS.log. The extra period between the date and time variables is not a typographical error, but rather an instance 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. The PATHMODE subparameter sets the file mode to 775 and the PATHOPTS subparameter OWRONLY opens the file for WRITE access. The sub-parameter OCREAT indicates that if the file does not already exist, create it.

You can modify the SYSPRINT DD card in either your Servant or Controller startup procedure. In addition, the SYSOUT DD card can be modified in the same way to redirect the SYSOUT output.

```
//*SYSPRINT DD SYSOUT=*,SPIN=UNALLOC,FREE=CLOSE
//SYSPRINT DD PATHMODE=(SIRWXU,SIRWXG,SIROTH),
// PATHOPTS=(OWRONLY,OCREAT),
// PATH='/myDir/myServer/was.log.d&LYMMDD..t&LHHMSS'
```

Note: If you try to direct the output for multiple streams to the same file, such as setting both DEFAULTDD and HRDCPYDD variables, the allocation for the HRDCPYDD file fails and output is sent to the default location (JOBLOG/SYSLOG).

Viewing JVM logs

The Java virtual machine (JVM) logs are written as plain text files.

About this task

IBM i Use either of two techniques to view the JVM logs for an application server:

- Use the administrative console, which also supports viewing the JVM logs from a remote machine.
- Use a text editor on the machine where the logs are stored.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log, SystemErr.log, trace.log, and activity.log files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

IBM i

Procedure

1. View the JVM logs from the administrative console.
 - a. Start the administrative console.

- b. Click **Troubleshooting > Logs and Trace** in the console navigation tree. To view the logs for a particular server, click on the server name to select it, then click **JVM Logs**.
 - c. Select the runtime tab.
 - d. Click **View** corresponding to the log you want to view.
2. View the JVM logs from the machine where they are stored.
 - a. Go to the machine where the logs are stored.
 - b. Navigate to the `profile_root/logs/server_name` directory and select `SystemOut.log` or `SystemErr.log`.
 - c. Open the file in a text editor or drag and drop the file into an editing and viewing program.

JVM log interpretation

View the JVM log files to determine problems within application environments.

The JVM logs contain print data written by applications. The application can write this data directly in the form of `System.out.print()`, `System.err.print()`, or other method calls. The application can also write data indirectly by calling a JVM function, such as an `Exception.printStackTrace()`. In addition, the `System.out` JVM log contains system messages written by the WebSphere Application Server.

Note: You can configure the thread ID that is used in the log and trace files of WebSphere Application Server Version 8.5 as either the hexadecimal representation of the thread ID from `java.util.logging.LogRecord`, or the hexadecimal representation of the thread ID from `java.lang.Thread`. Set the Java system property, `com.ibm.websphere.logging.useJULThreadID`, to `true` to have the thread ID match the `java.util.logging.LogRecord` thread ID. Set the system property to `false` to have the thread ID match the `java.lang.Thread` thread ID. If the system property is not specified, the thread ID matches the `java.lang.Thread` thread ID.

IBM i You can format application data to look like WebSphere Application Server system messages by using the Installed Application Output field of the JVM Logs properties panel, or as plain text with no additional formatting. WebSphere Application Server system messages are always formatted. Depending on how the JVM log is configured, formatted messages can be written to the JVM logs in either basic or advanced format.

IBM i

Message formats

Formatted messages are written to the JVM logs in one of two formats:

Basic Format

The format used in earlier versions of WebSphere Application Server.

Advanced Format

Extends the basic format by adding information about an event, when possible.

IBM i

Basic and advanced format fields

Basic and Advanced Formats use many of the same fields and formatting techniques. The various fields that may be found in these formats follow:

TimeStamp

The timestamp is formatted using the locale of the process where it is formatted. It includes a fully qualified date (for example `YYMMDD`), 24 hour time with millisecond precision and a time zone.

ThreadId

An 8 character hexadecimal value generated from the hash code of the thread that issued the message.

ThreadName

The name of the Java thread that issued the message or trace event.

ShortName

The abbreviated name of the logging component that issued the message or trace event. This is typically the class name for WebSphere Application Server internal components, but can be some other identifier for user applications.

LongName

The full name of the logging component that issued the message or trace event. This is typically the fully qualified class name for WebSphere Application Server internal components, but can be some other identifier for user applications.

EventType

A one character field that indicates the type of the message or trace event. Message types are in upper case. Possible values include:

- F** A Fatal message.
- E** An Error message.
- W** A Warning message.
- A** An Audit message.
- I** An Informational message.
- C** An Configuration message.
- D** A Detail message.
- O** A message that was written directly to System.out by the user application or internal components.
- R** A message that was written directly to System.err by the user application or internal components.
- Z** A placeholder to indicate the type was not recognized.

ClassName

The class that issued the message or trace event.

MethodName

The method that issued the message or trace event.

Organization

The organization that owns the application that issued the message or trace event.

Product

The product that issued the message or trace event.

Component

The component within the product that issued the message or trace event.

IBM i

Basic format

Message events displayed in basic format use the following format. The notation <name> indicates mandatory fields that will always appear in the basic format message. The notation [name] indicates optional or conditional fields that will be included if they can be determined.

```
<timestamp><threadId><shortName><eventType>[className][methodName]<message>
```

IBM i

Advanced format

Message events displayed in advanced format use the following format. The notation <name> is used to indicate mandatory fields that will always appear in the advanced format for message entries. The notation [name] is used to indicate optional or conditional fields that will be included if they can be determined.

```
<timestamp><threadId><eventType><UOW><source=longName>[className]
[methodName]<Organization><Product><Component>
[thread=threadName]<message>
```

Configuring the JVM logs

Use the administrative console to configure the JVM logs for an application server.

About this task

To log events or information from a running JVM, you can use the administrative console to configure the settings you need for each server. Configuration changes for the JVM logs that are made to a running application server are not applied until the application server is restarted.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Procedure

1. Start the administrative console
2. Click **Troubleshooting > Logs and Trace**, then click **server > JVM Logs**.
3. Select the Configuration tab.
4. Scroll through the panel to display the attributes for the stream to configure.
5. Change the appropriate configuration attributes and click **Apply**.
6. Save your configuration changes.

Java virtual machine (JVM) log settings

Use this page to view and modify the settings for the Java virtual machine (JVM) `System.out` and `System.err` logs.

Note: You can only access this page when the server is configured to use basic log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and Trace > server name > JVM Logs**.

View and modify the settings for the Java Virtual Machine (JVM) `System.out` and `System.err` logs for this managed process. The JVM logs are created by redirecting the `System.out` and `System.err` streams of the JVM to independent log files. The `System.out` log is used to monitor the health of the running application server. The `System.err` log contains exception stack trace information that is useful when performing problem analysis. There is one set of JVM logs for each application server and all of its applications. JVM logs are also created for the deployment manager and each node manager. Changes on the Configuration panel will apply when the server is restarted. Changes on the Runtime panel will apply immediately.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

File Name

Specifies the name of one of the log file described on this page.

The first file name field specifies the name of the System.out log. The second file name field specifies the name of the System.err file.

Note: This file cannot be shared between servers. The server creates a file with a .owner extension to help detect when two or more servers happen to be trying to use the same file.

Press the **View** button on the Runtime tab to view the contents of a selected log file.

The file name specified for the System.out log or the System.err log must have one of the following values:

filename

The name of a file in the file system. It is recommended that you use a fully qualified file name. If the file name is not fully qualified, it is considered to be relative to the current working directory for the server. Each stream must be configured with a dedicated file. For example, you cannot redirect both System.out and System.err to the same physical file.

If the directory containing the file already exists, the user ID under which the server is running requires read/write access to the directory. If the directory does not exist, it will be created with the proper permissions. The user id under which the server is running must have authority to create the directory.

console

This is a special file name used to redirect the stream to the corresponding process stream. If this value is specified for System.out, the file is redirected to stdout. If this value is specified for System.err, the file is redirected to stderr.

none Discards all data written to the stream. Specifying **none** is equivalent to redirecting the stream to dev/null on an operating system such as AIX® of Linux.

The default path for *filename* is the value of the variable SERVER_LOG_ROOT. To see the value of the SERVER_LOG_ROOT variable:

1. On the administrative console, select **Environment > WebSphere Variables**
2. Click on the **Server** radio button, and then click **Apply**. The value of the *SERVER_LOG_ROOT* variable appears in the resulting list.

To change the value of *SERVER_LOG_ROOT*:

1. Select *SERVER_LOG_ROOT*
2. Enter a new path in the Value field
3. Click Apply
4. Save the configuration. You will have to restart the server for the change to take effect.

You can also change the location and name of the `${SERVER_LOG_ROOT}/SystemOut.log` and `${SERVER_LOG_ROOT}/SystemErr.log` files to any other absolute path and filename (for example, `/tmp/myLogfile.log`).

File formatting

Specifies the format to use in saving the System.out file.

Log file rotation

Use this set of configuration attributes to configure the System.out or System.err log file to be self-managing.

A self-managing log file writes messages to a file until reaching either the time or size criterion. At the specified time or when the file reaches the specified size, logging temporarily suspends while the log file rolls over, which involves closing and renaming the saved file. The new saved file name is based on the original name of the file plus a timestamp qualifier that indicates when the renaming occurs. Once the

renaming completes, a new, empty log file with the original name reopens and logging resumes. All messages remain after the log file rollover, although a single message can split across the saved and the current file.

You can only configure a log to be self-managing if the corresponding stream is redirected to a file.

File Size

Click this attribute for the log file to manage itself based on its file size. Automatic roll over occurs when the file reaches the specified size you specify in the maximum size field.

Maximum Size

Specify the maximum size of the file in megabytes. When the file reaches this size, it rolls over.

This attribute is only valid if you click File size.

Time Click this attribute for the log file to manage itself based on the time of day. At the time specified in the start time field, the file rolls over.

Start Time

Specify the hour of the day, from 1 to 24, when the periodic rollover algorithm starts for the first time after an Application Server restart. The algorithm loads at Application Server startup. Once started at the (start time field) hour, the rollover algorithm rolls the file every (repeat time field) hours. This rollover pattern continues without adjustment until the Application Server stops.

Note: The rollover always occurs at the beginning of the specified hour of the day. The first hour of the day, which starts at 00:00:00 (midnight), is hour 1 and the last hour of the day, which starts at 23:00:00, is hour 24. Therefore, if you want log files to roll over at midnight, set the start time to 1.

Repeat time

Specifies the number of hours after which the log file rolls over. Valid values range from 1 to 24.

Configure a log file to roll over by time, by size, or by time and size. Click **File Size** and **Time** to roll the file at the first matching criterion. For example, if the repeat time field is 5 hours and the maximum file size is 2 MB, the file rolls every 5 hours, unless it reaches 2 MB before the interval elapses. After the size rollover, the file continues to roll at each interval.

Maximum Number of Historical Log Files

Specifies the number of historical (rolled) files to keep. The stream writes to the current file until it rolls. At rollover, the current file closes and is saved as a new name consisting of the current name plus the rollover timestamp. The stream then reopens a new file with the original name to continue writing. The number of historical files grows from zero to the value of the maximum number of historical files field. The next rollover deletes the oldest historical file.

Installed Application Output

Specifies whether System.out or System.err print statements issued from application code are logged and formatted.

Show application print statements

Click this field to show messages that applications write to the stream using **print** and **println** stream methods. WebSphere Application Server system messages always appear.

Format print statements

Click this field to format application print statement like WebSphere Application Server system messages.

Monitoring application logging using JMX notifications

Java developers can create programs to monitor application server logs using JMX notifications.

About this task

The most common log message listeners are written in Java, and connect to the deployment manager or an application server using SOAP. Use this topic to build a Java client that listens for log events.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Note: Be careful when adding listeners to servers with high logging volume as JMX notifications can slow down your server.

Procedure

1. Import the necessary packages. You will typically need the following import statements at the beginning of your Java program:

```
import javax.management.Notification;
import javax.management.NotificationListener;
import javax.management.ObjectName;
import javax.management.InstanceNotFoundException;
import javax.management.MalformedObjectNameException;

import com.ibm.websphere.management.AdminClient;
import com.ibm.websphere.management.AdminClientFactory;
import com.ibm.websphere.management.exception.ConnectorException;
```

Additionally, to handle the messages, and the types returned from the calls in subsequent steps you will need the following import statements.

```
import java.util.Iterator;
import java.util.Properties;
import java.util.Set;
import com.ibm.websphere.ras.RasMessage;
```

2. Create a Java class that implements the `NotificationListener` interface.
3. Implement the `handleNotification` method. The following example is a sample that writes the message text to the Java console:

```
public void handleNotification(Notification notification, Object handback) {
    RasMessage rasMessage = (RasMessage)notification.getUserData();
    System.out.println("Localized message: " + rasMessage.getLocalizedMessage(null));
}
```

4. Connect to the SOAP port of the server whose JMX MBeans you want to monitor. The following code creates a SOAP-connected `AdminClient` object with a specified host and a specified port:

```
AdminClient adminClient = null ;
String hostName = "someHostName";
String soapPort = "8880";

Properties connectProps = new Properties();
connectProps.setProperty(AdminClient.CONNECTOR_TYPE, "SOAP");
connectProps.setProperty(AdminClient.CONNECTOR_HOST, hostName);
connectProps.setProperty(AdminClient.CONNECTOR_PORT, soapPort);

try {
    adminClient = AdminClientFactory.createAdminClient(connectProps);
} catch (ConnectorException e) {
    // error handling code
}
```

5. Retrieve the MBean object name for the `RasLoggingService` MBean. The following code retrieves the `RasLoggingService` MBean object name:

```
String queryString = "WebSphere:cell="+cellName+",node="+nodeName+",process="+serverName+",
type=RasLoggingService,*" ;
Set<ObjectName> objectMBeans = null;
try {
    ObjectName queryName = new ObjectName(queryString);
```

```

        objectMBeans = (Set<ObjectName>)adminClient.queryNames(queryName, null);
        if (objectMBeans.size() > 1) {
            // error handling code to deal with the case where we get more than one name returned.
        }
    } catch (MalformedObjectNameException e) {
        // error handling code
    } catch (ConnectorException e) {
        // error handling code
    }

    if (objectMBeans.isEmpty()) {
        // error handling code to deal with the case where the MBean is not found
    }

    Iterator<ObjectName> objectNames = objectMBeans.iterator() ;
    ObjectName objectName = objectNames.next() ;

```

6. Add the notification listener. This sample code adds a notification listener, waits for 60 seconds while it processes notifications, then removes the notification listener. A listener can stay connected as long as needed.

```

    try {
        adminClient.addNotificationListener(objectName, this, null, null);
        Thread.sleep(60 * 1000) ;
        adminClient.removeNotificationListener(objectName, this) ;
    } catch (InstanceNotFoundException e) {
        // error handling code
    } catch (ConnectorException e) {
        // error handling code
    } catch (Exception e) {
        // error handling code
    }
}

```

7. Add the necessary jar to your classpath. Add the admin client jar file to your classpath to be able to compile and run your code. The admin client jar file is in the <install_root>/runtimes directory.

Results

You have created a Java program that can listen to, and take actions as a result of log event notifications from an application server.

Process logs

WebSphere Application Server processes contain two output streams that are accessible to native code running in the process. These streams are the stdout and stderr streams. Native code, including Java virtual machines (JVM), might write data to these process streams. In addition, JVM provided System.out and System.err streams can be configured to write their data to these streams also.

IBM i By default, the stdout and stderr streams are redirected to log files at application server startup, which contain text written to the stdout and stderr streams by native modules (*SRVPGMs, .dlls, .exes, UNIX libraries, and other modules). By default, these files are stored as *profile_root/logs/server_name/native_stderr.log* and *profile_root/logs/native_stdout.log*.

Important: If you turn on the debugging service or verbose Java logging, the size of the *native_stderr.log* and *native_stdout.log* files increase in size with each restart of the application server. This situation can cause application server crashes especially on Linux operating systems where you might have a 2 gigabyte maximum file size. To prevent this problem, consider redirecting the Java exceptions and verbose Java garbage collection to other files.

Configuring the service log

The settings for service logs are typically shared for all servers, but you can configure a separate service log for each server process by overriding the configuration values at the server level.

About this task


The configuration values for the service log are inherited by each server process from the node configuration, but under certain circumstances you might wish to configure the service logs differently for individual servers. You can use the administrative console to change the service log settings from the server level configuration panels.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

In WebSphere Application Server Version 7 and earlier, the service log is enabled by default. In WebSphere Application Server Version 8x, however, the service log is disabled by default. Configure your servers to use HPEL log and trace mode and use the HPEL API if you need to be able to merge log file content from multiple servers. Use the HPEL log, trace mode and the HPEL LogViewer command if you need to be able to render log content in Common Base Event XML format.

Procedure

1. Start the administrative console.
2. Click **Troubleshooting > Logs and Trace > *server_name* > IBM Service Logs**.
3. Select the **Enable** box to enable the service log, or clear the check box to disable the log.
4. Set the name for the service log.

 The default name is `profile_root/logs/activity.log`. If the name is changed, the run time requires write access to the new file, and the file must use the `.log` extension.

5. Set the maximum file size.
Specifies the number of megabytes to which the file can grow. When the file reaches this size, it wraps, replacing the oldest data with the newest data.
6. Save the configuration.
7. Restart the server to apply the configuration changes.

IBM service log settings

Use this page to configure the IBM service log, also known as the activity log.

Note: You can only access this page when the server is configured to use basic log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and Trace > *server name* > IBM Service Logs**.

Use this page to configure the IBM service log, also known as the activity log. The IBM service log contains both the WebSphere Application Server messages that are written to the `System.out` stream and some special messages that contain extended service information that can be important when analyzing problems. There is one service log for all WebSphere Application Server Java virtual machines (JVMs) on a node, including all application servers. The IBM Service log is maintained in a binary format. Use the Log and Trace Analyzer or Showlog tool to view the IBM service log.

Enable service log

Specifies creation of a log file by the IBM Service log.

File Name

Specifies the name of the file used by the IBM Service log.

Maximum File Size

Specifies the maximum size in megabytes of the service log file. The default value is 2 megabytes.

When this size is reached, the service log wraps in place. Note that the service log does not roll over to a new log file like the JVM logs.

Enable Correlation ID

Specifies the generation of a correlation ID that is logged with each message.

You can use the correlation ID to correlate activity to a particular client request.

Viewing the service log

Service logs are logs written in a binary format. You cannot view a service log directly using a text editor. You should never directly edit the service log, as doing so will corrupt the log.

Before you begin

IBM i To move a service log from one machine to another, you must use a mechanism like FTP, which supports binary file transfer. Use the Showlog tool to convert the contents of the service log to a text format that you can then write to a file or dump to the command shell window.

About this task

Run the showlog script to view the contents of the service log as described in the following procedure.

Procedure

1. Open a shell window on the machine where the service log resides.
2. Change the directory to *app_server_root/bin* where *app_server_root* is the fully qualified path where the WebSphere Application Server product is installed.
3. Run the showlog script.

IBM i `showlog`

4. Run the following showlog script with no parameters to display usage instructions.

IBM i `showlog`

5. Format and write the service log contents to a file.

`showlog service_log_filename output_filename`

If the service log is not in the default location, you must fully qualify the *service_log_filename*

Chapter 11. Working with trace


Use trace to obtain detailed information about running the WebSphere Application Server components, including application servers, clients, and other processes in the environment.

About this task

Trace files show the time and sequence of methods called by WebSphere Application Server base classes, and you can use these files to pinpoint the failure. Collecting a trace is often requested by IBM technical support personnel. If you are not familiar with the internal structure of WebSphere Application Server, the trace output might not be meaningful to you.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Procedure

1. Configure an output destination to which trace data is sent.
2. Enable trace for the appropriate WebSphere Application Server or application components.
3. Run the application or operation to generate the trace data.
4.  Analyze the trace data or forward it to the appropriate organization for analysis.

Results

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.

Enabling trace on client and stand-alone applications

When stand-alone client applications (such as Java applications which access enterprise beans hosted in WebSphere Application Server) have problems interacting with WebSphere Application Server, it might be useful to enable tracing for the application. Enabling trace for client programs will cause the WebSphere Application Server classes used by those applications, such as naming-service client classes, to generate trace information.

About this task

A common troubleshooting technique is to enable tracing on both the application server and client applications, and match records according to timestamp to try to understand where a problem is occurring.

Procedure

1. To enable trace for the WebSphere Application Server classes in a client application, add the system properties shown in the following example to the startup script or command of the client application. The location of the output and the classes and detail included in the trace follow the same rules as for adding trace to WebSphere Application Servers. For example, trace the stand-alone client application program named `com.ibm.sample.MyClientProgram`, enter the following command:

```
java -DtraceSettingsFile=MyTraceSettings.properties
-Djava.util.logging.manager=com.ibm.ws.bootstrap.WsLogManager
-Djava.util.logging.configureByServer=true com.ibm.samples.MyClientProgram
```

IBM i The file identified by *file name* must be a properties file placed in the class path of the application client or stand-alone process. You must create a trace properties file by copying the `app_server_root/properties/TraceSettings.properties` file to the same directory as your client application Java archive (JAR) file

You cannot use the `-DtraceSettingsFile=TraceSettings.properties` property to enable tracing of the ORB component for thin clients. ORB tracing output for thin clients can be directed by setting `com.ibm.CORBA.Debug.Output = debugOutputFilename` parameter in the command line.

The `java.util.logging.manager` and `java.util.logging.configureByServer` system properties configure Java logging to use a WebSphere Application Server-specific LogManager class and to use the configuration from the file specified by the `traceSettingsFile` property. The default Java Logging properties file, located in the Java SE Runtime Environment 6 (JRE6), will not be applied.

- IBM i** You can configure the `MyTraceSettings.properties` file to send trace output to a file using the `traceFileName` property. Specify one of two options:
 - The fully qualified name of an output file. For example, `traceFileName=c:\\MyTraceFile.log`. You must specify this property to generate visible output.
 - `stdout`. When specified, output is written to `System.out`.
- You can also specify a trace string for writing messages with the `Trace String` property. Specify a startup trace specification similar to that available on the server. For your convenience, you can enter multiple individual trace strings into the trace settings file, one trace string per line.

Results

Here are the results of using each optional property setting:

- Specify a valid setting for the `traceFileName` property without a trace string to write messages to the specified file or `System.out` only.
- Specify a trace string without a `traceFileName` property value to generate no output.
- Specify both a valid `traceFileName` property and a trace string to write both message and trace entries to the location specified in the `traceFileName` property.

Enabling trace at server startup

Use the administrative console to enable tracing at a server's startup. You can use trace to assist you in monitoring system performance and diagnosing problems.



About this task

The diagnostic trace configuration settings for a server process determines the initial trace state for a server process. The configuration settings are read at server startup and used to configure the trace service. You can also change many of the trace service properties or settings while the server process is running.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Procedure

1. Start the administrative console.

2.  Click **Troubleshooting > Logs and trace** in the console navigation tree, then click **Server > Diagnostic Trace**.
3. Click **Configuration**.
4.  Do not select the **None** check box. If this option is selected, the trace data is not logged or recorded anywhere. All other handlers (including handlers registered by applications) still have an opportunity to process these traces.
5. Select whether to direct trace output to either a file or an in-memory circular buffer.

Note: Different components can produce different amounts of trace output per entry. Naming and security tracing, for example, produces a much higher trace output than web container tracing. Consider the type of data being collected when you configure your memory allocation and output settings.
6. If the in-memory circular buffer is selected for the trace output set the size of the buffer, specified in thousands of entries. This is the maximum number of entries that will be retained in the buffer at any given time.
7. If a file is selected for trace output, set the maximum size in megabytes to which the file should be allowed to grow. When the file reaches this size, the existing file will be closed, renamed, and a new file with the original name reopened. The new name of the file will be based upon the original name with a timestamp qualifier added to the name. In addition, specify the number of history files to keep.
8. Select the desired format for the generated trace.
9. Save the changed configuration.
10. To enter a trace string to set the trace specification to the desired state:
 - a. Click **Troubleshooting > Logs and trace** in the console navigation tree.
 - b. Select a server name.
 - c. Click **Change Log Level Details**.
 - d. If **A11 Components** has been enabled, you might want to turn it off, and then enable specific components.
 - e. Click a component or group name. For more information see the page on log level settings. If the selected server is not running, you will not be able to see individual component in graphic mode.
 - f. Enter a trace string in the trace string box.
 - g. Select **Apply**, then **OK**.
11. Allow enough time for the nodes to synchronize, and then start the server.

Enabling trace on a running server


Use the administrative console to enable tracing on a running server. You can use trace to assist you in monitoring system performance and diagnosing problems.

About this task

You can modify the trace service state that determines which components are being actively traced for a running server by using the following procedure.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Procedure

1. Start the administrative console.
2.  Click **Troubleshooting > Logs and Trace** in the console navigation tree, then click **server > Diagnostic Trace**.
3. Select the **Runtime** tab.
4. Select the **Save runtime changes to configuration as well** check box if you want to write your changes back to the server configuration.
5. Change the existing trace state by changing the trace specification to the desired state. Refer to the following topics for specific information on how to change your trace specifications:
 - “Diagnostic trace service settings” describes how you can examine the current diagnostic trace settings and to change where the trace is written (memory buffer or file). If the trace is to be recorded in a memory buffer, you can specify the file to which the memory buffer is to be dumped. You can view the file for trace information.
 - “Log and trace settings” describes how you can view and configure logging and trace settings for the server.
 - “Managing the application server trace service” describes how you can manage the trace service for a server process while the server is stopped and while it is running. You can specify which components to trace, where to send trace output, the characteristics of the trace output device, and which format to generate trace output in. You can select where the trace is written (memory buffer or file).

On an application server, trace output can be directed either to a file or to an in-memory circular buffer. If trace output is directed to the in-memory circular buffer, it must be dumped to a file before it can be viewed.

On an application client or stand-alone process, trace output can be directed either to a file or to the process console window.
6. Configure the trace output if a change from the existing one is desired.
7. Click **Apply**.

Managing the application server trace service

You can manage the trace service for a server process while the server is stopped and while it is running. You can specify which components to trace, where to send trace output, the characteristics of the trace output device, and which format to generate trace output in.

About this task

Modify the trace settings to help with diagnosing problems or tuning performance within certain applications.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log`, `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Procedure

1. Start the administrative console.
2. Click **Troubleshooting > Logs and Trace > *server_name***.
3. Click the **Diagnostic Trace** link.

4. **IBM i** On the Configuration tab, do not select the **None** option. If this option is selected, the trace data is not logged or recorded anywhere. All other handlers (including handlers registered by applications) still have an opportunity to process these traces.
5. **IBM i** On the Runtime tab, select either the **Memory Buffer** or **File Trace Output** option.
6. Specify the appropriate values for your configuration for either the Memory Buffer or File Trace Output option.
7. Click **Apply**.
8. Click **Troubleshooting > Logs and Trace > *server_name***.
9. Click the **Change Log Detail Levels** link.
10. Under Additional properties, click **Change Log Detail Levels**.
11. On the Runtime tab, change the existing trace state by changing the trace specification to the desired state.
12. Click **Apply**.

Diagnostic trace service settings

Use this page to configure diagnostic trace service settings.

Note: You can only access this page when the server is configured to use basic log and trace mode.

To view this administrative console page, click **Troubleshooting > Logs and trace > *server_name* > Diagnostic trace**.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log, SystemErr.log, trace.log, and activity.log files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Trace Output

Specifies where trace output should be written. The trace output can be written directly to an output file, or stored in memory and written to a file on demand using the Dump button found on the runtime page.

Different components can produce different amounts of trace output per entry. Naming and security tracing, for example, produces a much higher trace output than web container tracing. Consider the type of data being collected when you configure your memory allocation and output settings.

None If this option is selected, the trace data is not logged or recorded anywhere. All other handlers (including handlers registered by applications) still have an opportunity to process these traces.

Memory Buffer

Specifies that the trace output should be written to an in-memory circular buffer. If you select this option you must specify the following parameters:

Maximum Buffer Size

Specifies the number of entries, in thousands, that can be cached in the buffer. When this number is exceeded, older entries are overwritten by new entries.

File Specifies to write the trace output to a self-managing log file. The self-managing log file writes messages to the file until the specified maximum file size is reached. When the file reaches the specified size, logging is temporarily suspended and the log file is closed and renamed. The new name is based on the original name of the file, plus a timestamp qualifier that indicates when the renaming occurred. Once the renaming is complete, a new, empty log file with the original name is

reopened, and logging resumes. No messages are lost as a result of the rollover, although a single message may be split across the two files. If you select this option you must specify the following parameters:

Maximum File Size

Specifies the maximum size, in megabytes, to which the output file is allowed to grow. This attribute is only valid if the File Size attribute is selected. When the file reaches this size, it is rolled over as previously described.

Maximum Number of Historical Files

Specifies the maximum number of rolled over files to keep.

File Name

Specifies the name of the file to which the trace output is written.

Note: This file cannot be shared between servers. The server creates a file with a .owner extension to help detect when two or more servers happen to be trying to use the same file.

Trace Output Format

Specifies the format of the trace output.

You can specify one of three levels for trace output:

Basic (Compatible)

Preserves only basic trace information. Select this option to minimize the amount of space taken up by the trace output.

Advanced

Preserves more specific trace information. Select this option to see detailed trace information for use in troubleshooting and problem determination.

Log analyzer trace format

Preserves trace information in the same format as produced by Showlog tool.

Runtime tab

Save runtime changes to configuration

Save runtime changes made on the runtime tab to the trace configuration as well. Select this box to copy runtime trace changes to the trace configuration settings as well. Saving these changes to the trace configuration will cause the changes to persist even if the application is restarted.

Trace Output

Specifies where trace output should be written. The trace output can be written directly to an output file, or stored in memory and written to a file on demand using the Dump button found on the run-time page.

None If this option is selected, the trace data is not logged or recorded anywhere. All other handlers (including handlers registered by applications) still have an opportunity to process these traces.

Memory Buffer

Specifies that the trace output should be written to an in-memory circular buffer. If you select this option you must specify the following parameters:

Maximum Buffer Size

Specifies the number of entries, in thousands, that can be cached in the buffer. When this number is exceeded, older entries are overwritten by new entries.

Dump File Name

The name of the file to which the memory buffer will be written when it is dumped. This option is only available from the Runtime tab.

File Specifies to write the trace output to a self-managing log file. The self-managing log file writes messages to the file until the specified maximum file size is reached. When the file reaches the specified size, logging is temporarily suspended and the log file is closed and renamed. The new name is based on the original name of the file, plus a timestamp qualifier that indicates when the renaming occurred. Once the renaming is complete, a new, empty log file with the original name is reopened, and logging resumes. No messages are lost as a result of the rollover, although a single message may be split across the two files. If you select this option you must specify the following parameters:

Maximum File Size

Specifies the maximum size, in megabytes, to which the output file is allowed to grow. This attribute is only valid if the File Size attribute is selected. When the file reaches this size, it is rolled over as previously described.

Maximum Number of Historical Files

Specifies the maximum number of rolled over files to keep.

File Name

Specifies the name of the file to which the trace output is written.

View View the file that is specified by the **File Name** parameter. This does not apply your configuration.

Note: Do not specify the same file name for more than one managed server. The server creates a file with a .owner extension to track which server can use a particular file name for output.

Select a server to configure logging and tracing

Use this page to select the server for which you want to configure logging and trace settings.

This page lists application servers in the cell and the nodes holding the application servers. The status indicates whether a server is running, stopped, or encountering problems.

When you select an application server, a page is displayed that will allow you to choose which log or trace task to configure for that application server.

To view this administrative console page, click **Troubleshooting > Logs and Trace**.

Server

Specifies the logical name of the server.

Host name

Specifies the name of the host for the application server.

Version

Specifies the version for the application server.

Type

Specifies the type of application server.

Status

Indicates whether the application server is started or stopped. (WebSphere Application Server, Network Deployment only)

Note that if the status is *Unavailable*, the node agent is not running in that node, and you must restart the node agent before you can start the server.

Log and trace settings

Use this page to view and configure logging and trace settings for the server.

Note: You can only access this page when the server is configured to use basic log and trace mode.

To view this administrative console page, click **Troubleshooting** > **Logs and trace** > *server_name* .

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using `SystemOut.log` , `SystemErr.log`, `trace.log`, and `activity.log` files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Switch to HPEL Mode

IBM i

Use the Switch to HPEL Mode button to change the log and trace mode for the server to HPEL.

Note: Switching the server to HPEL log and trace mode requires a server restart.

Diagnostic Trace

IBM i

The diagnostic trace configuration settings for a server process determine the initial trace state for a server process. The configuration settings are read at server startup and used to configure the trace service. You can also change many of the trace service properties or settings while the server process is running.

Java virtual machine (JVM) Logs

IBM i

The JVM logs are created by redirecting the `System.out` and `System.err` streams of the JVM to independent log files. WebSphere Application Server writes formatted messages to the `System.out` stream. In addition, applications and other code can write to these streams using the `print()` and `println()` methods defined by the streams.

Process Logs

IBM i

WebSphere Application Server processes contain two output streams that are accessible to native code running in the process. These streams are the `stdout` and `stderr` streams. Native code, including Java virtual machines (JVM), might write data to these process streams. In addition, JVM provided `System.out` and `System.err` streams can be configured to write their data to these streams also.

IBM Service Logs

IBM i

The IBM service log contains both the WebSphere Application Server messages that are written to the System.out stream and some special messages that contain extended service information that is normally not of interest, but can be important when analyzing problems. There is one service log for all WebSphere Application Server JVMs on a node, including all application servers.

Change Log Level Details

Enter a log detail level that specifies the components, packages, or groups to trace. The log detail level string must conform to the specific grammar described in this topic. You can enter the log detail level string directly, or generate it using the graphical trace interface.

NCSA access and HTTP error logging

The NCSA access and HTTP error logging page enables you to configure the log settings for your HTTP server.

Chapter 12. Troubleshooting class loaders

Class loaders find and load class files. For a deployed application to run properly, the class loaders that affect the application and its modules must be configured so that the application can find the files and resources that it needs. Diagnosing problems with class loaders can be complicated and time-consuming. To diagnose and fix the problems more quickly, use the administrative console class loader viewer to examine class loaders and the classes loaded by each class loader.

Before you begin

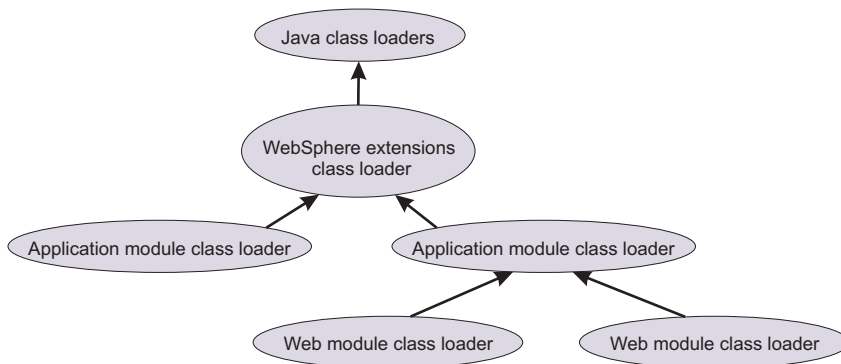
This topic assumes that you have installed an application on a server supported by the product and you want to examine class loaders used by the application or its modules. The modules can be web modules (.war files) or enterprise bean (EJB) modules (.jar files). The class loader viewer enables you to examine class loaders in a runtime environment.

This topic also assumes that you have enabled the class loader viewer service. Click **Servers > Server Types > WebSphere application servers > server_name > Class loader viewer service**, enable the service and restart the server.

About this task

The runtime environment of WebSphere Application Server uses the following class loaders to find and load new classes for an application in the following order:

1. The bootstrap, extensions, and CLASSPATH class loaders created by the Java virtual machine
2. A WebSphere extensions class loader
3. One or more application module class loaders that load elements of enterprise applications running in the server
4. Zero or more web module class loaders



Each class loader is a child of the previous class loader. That is, the application module class loaders are children of the WebSphere extensions class loader, which is a child of the CLASSPATH Java class loader. Whenever a class needs to be loaded, the class loader usually delegates the request to its parent class loader. If none of the parent class loaders can find the class, the original class loader attempts to load the class. Requests can only go to a parent class loader; they cannot go to a child class loader. After a class is loaded by a class loader, any new classes that it tries to load reuse the same class loader or go up the precedence list until the class is found.

If the class loaders that load the artifacts of an application are not configured properly, the Java virtual machine (JVM) might throw a class loading exception when starting or running that application. “Class

loading exceptions” on page 163 describes the types of exceptions caused by improperly configured class loaders and suggests ways to use the class loader viewer to correct configurations of class loaders. The types of exceptions include:

- ClassCastException
- ClassNotFoundException
- NoClassDefFoundException
- UnsatisfiedLinkError

Use the class loader viewer to examine class loaders and correct problems with application or class loader configurations.

Procedure

- Examine a tree view that lists all installed applications and their modules. The modules can be web modules (.war files) or EJB modules (.jar files).

Click **Troubleshooting** > **Class loader viewer** to access the Enterprise applications topology page.

- Examine the class loader delegation hierarchy.

On the Enterprise applications topology page, select a module to access the Class loader viewer page. The page lists the class loaders visible to web and EJB modules in an installed enterprise application. This page helps you to determine which class loaders loaded files of a module and to diagnose problems with class loaders.

The delegation hierarchy is determined by the class loader delegation mode, or *class loader order*, specified for an application or web module. The value can be either `Classes loaded with parent class loader first` or `Classes loaded with local class loader first (parent last)`. Refer to the *Configure class loaders* step for more information.

- Export information on class loaders.
 1. On the Class loader viewer page, click **Export**.
 2. Select to open a browser or editor on the class loader information or to save the information to disk in XML format.
 3. Click **OK**, and specify any additional information requested by the system.
- Display information about class loaders visible to the module in an HTML table format.

On the Class loader viewer page, click **Table View**. The Table View page displays the following information:

Table 32. Table View page. Information available on class loader attributes.

Class loader attribute	Description
Delegation	Indicates whether the class loader delegates the loading of the module to its parent class loader. A value of <code>true</code> implies that the class loader of the parent application is being used (<code>Classes loaded with parent class loader first</code>). A value of <code>false</code> implies that the module class loader is being used (<code>Classes loaded with local class loader first (parent last)</code>). Refer to the <i>Configure class loaders</i> step for more information.
Classpath	Lists the paths over which the class loader searches for classes and resources.
Classes	Lists the names of classes loaded in the JVM by this class loader.

The **Table View** option does not return a value when out-of-memory errors are generated. The out-of-memory errors might be related to a memory leak. To examine information about class loaders in a table, resolve the out-of-memory problem, and then click **Table View** again.

- Search class loaders.

On the Class loader viewer page, click **Search** to access the Search page, on which you can search class loaders for the following:

- Specific strings
- Specific .jar files
- The names of files in a specific directory

- The names of files loaded by a specific class loader

The search is case-sensitive. “Class loading exceptions” describes several uses of the Search page.

- Configure class loaders. You can configure class loaders for the following:
 - All applications installed on a specific server.
 - A specific application
 - A specific web module

Note: For detailed information about server, application, and web class loaders, see the chapter on class loading in the *Developing and deploying applications* PDF book.

Class loader configuration determines which class loader loads the classes and resource files for an application or web module. Application and WAR module class loader configuration settings include **Class loader order** and **WAR class loader policy**.

A **Class loader order** value can be either `Classes loaded with parent class loader first` or `Classes loaded with local class loader first (parent last)`. The default is `Classes loaded with parent class loader first`. A class loader with the `Classes loaded with parent class loader first` mode delegates loading a class or resource to its immediate parent class loader before searching its classpath.

When troubleshooting class loading problems, you might need to override classes visible to a parent class loader. To override such classes with those specific to an application, set the **Class loader order** to `Classes loaded with local class loader first (parent last)` on the class loader that contains the application classes on its classpath. An application can override classes visible to a parent class loader, but doing so can result in a `ClassCastException` or `UnsatisfiedLinkError` if there is a mixed use of overridden classes and non-overridden classes.

For example, under default class loader policies, a web module has its own Web module (WAR) class loader to load its artifacts, which are typically in the `WEB-INF/classes` and `WEB-INF/lib` directories. An application module class loader is the immediate parent of this WAR class loader. To ensure that the web module class loader searches these paths for a particular class or resource first, before delegating the load operation to the application module class loader, set the **Class loader order** of the web module to `Classes loaded with local class loader first (parent last)`.

Class loader policies determine the structure of the application and WAR module class loaders. Under the default policies, every running application EAR has its own application module class loader, and every web module has its own WAR module class loader. The default policies ensure Java EE compliance regarding visibility and isolation among application artifacts. Changing the default policies is not suggested when troubleshooting class loading problems.

What to do next

If you continue to have class loader problems, refer to “Class loading exceptions” and to the class loading chapter of the *Developing and deploying applications* PDF book.

Class loading exceptions

What kind of class-loading error do you see when you develop an application or start an installed application?

- “`ClassCastException`”
- “`ClassNotFoundException`” on page 165
- “`NoClassDefFoundException`” on page 166
- “`UnsatisfiedLinkError`” on page 166

ClassCastException

A class cast exception results when the following conditions exist and can be corrected by the following actions:

- The type of the source object is not an instance of the target class (type).

- The class loader that loaded the source object (class) is different from the class loader that loaded the target class.
- The application fails to perform or improperly performs a narrow operation.

The type of the source object is not an instance of the target class (type).

This is the typical class cast exception. You can diagnose whether the source object of a cast statement is not an instance of the target class (type) by examining the class signature of the source object class, then verifying that it does not contain the target class in its ancestry and the source object class is different than the target class. You can obtain class information by inserting a simple print statement in your code. For example:

```
System.out.println( source.getClass().getName() + ":" + target.getClass().getName() );
```

Or use a **javap** command. For example:

```
javap java.util.HashMap
Compiled from "HashMap.java"
public class java.util.HashMap extends java.util.AbstractMap
    implements java.util.Map,java.lang.Cloneable,java.io.Serializable {
```

The class loader that loaded the source object (class) is different from the class loader that loaded the target class.

Assuming that the type of the source object is an instance of the target class, a class cast exception occurs when the class loader that loaded the source object's class is different than the class loader that loaded the target class. This condition might occur when the target class is visible on the classpaths of more than one class loader in the WebSphere Application Server runtime environment. To correct this problem, use the Search and Search by class name console pages used to diagnose problems with class loaders:

1. Click **Troubleshooting** > **Class loader viewer** > *module_name* > **Search** to access the Search page.
2. For **Search type**, select **Class/Package**.
3. For **Search terms**, type the name of the class that is loaded by two class loaders.
4. Click **OK**. The Search by class name page is displayed, listing all class loaders that load the class.

If there is more than one class loader listed, then the target class was loaded by more than one class loader. Because the source object is an instance of the target class, the class loader that loaded the source object class is different from the class loader that loaded the target class.

5. Return to the Class loader viewer page and examine the classpath to determine why two different class loaders load the class.
6. Correct your code so that the class is visible only to the appropriate class loader.

The application fails to perform or improperly performs a narrow operation.

A class cast exception can occur because, when the application is resolving a remote enterprise bean (EJB) object, the application code does not perform a narrow operation as required. The application must perform a narrow operation after looking up a remote object. Examine the application and determine whether it looks up a remote object and, if so, the result of the lookup is submitted to a narrow method.

The narrow method must be invoked according to the EJB 2.0 programming model. In particular, the target class submitted to the narrow method must be the exact, most derived interface of the EJB. This also causes a class cast exception in the WebSphere Application Server runtime environment. Examine the application and determine whether the target class submitted to the narrow method is a super-interface of the EJB that is specified, not the exact EJB type; if so, modify the application to invoke narrow with the exact EJB interface.

Lastly, if a class cast exception occurs during a narrow operation, verify that the narrow method is being applied to the result of a remote EJB lookup, not to a local enterprise bean. A narrow is not used for local lookups. Examine the application or module deployment descriptor to ensure that the object being narrowed is not a local object.

ClassNotFoundException

A class not found exception results when the following conditions exist and can be corrected by the following actions:

- The class is not visible on the logical classpath of the context class loader.
- The application incorrectly uses a class loader API.
- A dependent class is not visible.

The class is not visible on the logical classpath of the context class loader.

The class not found is not in the logical class path of the class loader associated with the current thread. The logical classpath is the accumulation of all classpaths searched when a load operation is invoked on a class loader. To correct this problem, use the Search page to search by class name and by Java archive (JAR) name:

1. Click **Troubleshooting** > **Class loader viewer** > *module_name* > **Search** to access the Search page.
2. For **Search type**, select **Class/Package**.
3. For **Search terms**, type the name of the class that is not found.
4. Click **OK**. The Search by class name page is displayed, listing all class loaders that load the class.
5. Examine the page to see if the class exists in the list.
6. If the class is not in the list, return to the Search page. For **Search terms**, type the name of the .jar file for the class; for **Search type**, select **JAR/Directory**.
7. Click **OK**. The Search by Path page is displayed, listing all directories that hold the JAR file.

If the JAR file is not in the list, the class likely is not in the logical class path, not readable or an alternate class is already loaded. Move the class to a location that enables it to be loaded.

The application incorrectly uses a class loader API.

An application can obtain an instance of a class loader and call either the loadClass method on that class loader, or it can call Class.forName(*class_name*, *initialize*, *class_loader*) with that class loader. The application may be incorrectly using the class loader application programming interface (API). For example, the class name is incorrect, the class is not visible on the logical classpath of that class loader, or the wrong class loader was engaged.

To correct this problem, determine whether the class exists and whether the application is properly using the class loader API. Follow the steps in The class is not visible on the logical classpath of the context class loader to determine whether the class is loaded. If the class has not been loaded, attempt to correct the application and see if the class loads. If the class is in the class path with proper permission and is not being overridden by another factory class, examine the API used to load the class.

1. Click **Troubleshooting** > **Class loader viewer** > *module_name* > **Search** to access the class loader Search page.
2. For **Search type**, select **Class/Package**.
3. For **Search terms**, type the name of the class.
4. Click **OK**. The Search by class name page is displayed, listing all class loaders that load the class.
5. Examine the page to see if the class exists in the list.
6. If the class is in the list and a ClassNotFoundException exception was thrown, then the .jar file or class is not in the correct context or a wrong API call in the current context was used.

If the class is not in the list, return to the Search page and do the following:

- a. Search for the class that generated the exception; that is, the class calling Class.forName.
- b. See which class loader loads the class.
- c. Determine whether the class loader has access or can load the class not found by evaluating the class path of the class loader.

A dependent class is not visible.

When a class loader *clsldr* loads a class *cls*, the Java virtual machine (JVM) invokes *clsldr* to load the classes on which *cls* depends. Dependent classes must be visible on the logical classpath of *clsldr*, otherwise an exception occurs. This condition typically occurs when users make WebSphere Application Server classes visible to the JVM, or make application classes visible to the JVM or to the WebSphere extensions class loader. For example:

- Class A depends on Class B.
- Class A is visible to the WebSphere extensions class loader.
- Class B is visible on the local classpath of a WAR module class loader, not the WebSphere extensions class loader classpath.

When the JVM loads class A using the WebSphere extensions class loader, it then attempts to load Class B using the same class loader and ultimately creates a class not found exception.

To correct this problem:

1. Make the application-specific classes visible to the appropriate application class loader.
2. Search for the class not found (Class B).
3. If Class B is in the proper location, search for the class that loads the dependent class (Class A) in the Class loader viewer.
4. If the class is loaded and a `ClassNotFoundException` exception was thrown, then the `.jar` file or class is not in proper context or the wrong API call in the current context was used.

If no class was found, do the following:

- a. Search for the class that generated the exception; that is, the class calling `Class.forName`.
 - b. See which class loader loads the class.
 - c. Determine whether the class loader has access or can load the class not found by evaluating the class path of the class loader.
5. Ensure that the caller class (Class B) is visible to the JVM or WebSphere extensions class loader.

NoClassDefFoundException

A no class definition found exception results when the following conditions exist and can be corrected by the following actions:

The class is not in the logical class path.

Refer to “`ClassNotFoundException`” on page 165 for information.

The class cannot load.

There are various reasons for a class not loading. The reasons include: failure to load the dependent class, the dependent class has a bad format, or the version number of a class.

UnsatisfiedLinkError

A linkage error results when the following conditions exist and can be corrected by the following actions:

- A user action caused the error.
- The native library is already loaded.
- A dependent native library was used.

A user action caused the error.

Several user actions can result in a linkage error:

A library extension name is incorrect for the platform.

`System.loadLibrary` is passed an incorrect parameter.

The library is not visible.

As a best practice, use the JVM class loader to find or load native libraries. WebSphere Application Server prints the Java library path (`java.library.path`) when starting up. If the JVM class loader is intended to load the library, verify that the path containing the native

library file is in the Java library path. If not, append the path to the platform-specific native library environment variable or to the `java.library.path` system property of the server process definition.

In general, the Java virtual machine invokes `findLibrary()` on the class loader `xxx` that loads the class that calls `System.loadLibrary()`. If `xxx.findLibrary()` fails, the Java virtual machine attempts to find the library using the JVM class loader, which searches the JVM library path. If the library cannot be found, the Java virtual machine creates an `UnsatisfiedLinkError` exception.

Thus, if a WebSphere class loader is intended to find a native library `myNativeLib`, the library must be visible on the `nativeLibpath` of the class loader that loads the class that calls `System.loadLibrary(myNativeLib)`. This practice is necessary or desirable in the following situation:

- Shared libraries have a **Native library path** in their configuration. Because shared libraries enable the versioning of application-specific libraries, consider specifying the paths to any native libraries used by the shared library code in the shared library configuration.

Ensure that the correct WebSphere class loader loads the class that calls `System.loadLibrary()` and that the native library is visible on the **Native library path** setting.

The native library is already loaded.

This condition can result from either of the following errors:

User error

Check for multiple calls to `System.loadLibrary` and remove any extraneous calls.

Error when an application restarts

The JVM has a restriction that only one class loader can load a native library at a time. An error results when an application restarts before the garbage collector cleans up the class loader from the stopped application. When the class that loads the native library moves, all of the classes that depend on that native library and their dependencies also must move.

To correct this condition, move the loading of the native library to a class loader that does not reload:

1. Locate all application classes that load native libraries or have native methods.
2. Identify any dependent classes for the classes in step 1, such as logging packages.
3. Create a server-associated shared library or an isolated shared library.
4. Move the JAR files loaded for classes in steps 1 and 2 from the application to the shared library created in step 3.
5. Save your changes.
6. Redeploy the application and rerun the scenario.

For more information about invoking, creating, and managing shared libraries, read “Managing shared libraries” in the *Administering applications and their environment* PDF book.

Classes within server-scoped libraries are loaded once for each server lifecycle, ensuring that the native library required by the application is loaded once for each Java virtual machine, regardless of the application's life cycle.

A dependent native library was used.

Dependent native libraries must be found or loaded by the JVM class loader. That is, if a native library `NL` is dependent on another native library, `DNL`, the JVM class loader must find `DNL` on the Java library path. This is because the JVM runs native code when loading `NL`; when it encounters the dependency on `DNL`, the JVM native code can call only to the JVM class loader to resolve the dependency. A WebSphere class loader cannot load a dependent native library.

Modify the platform-specific environment variable defining the Java library path (LIBPATH) to include the path containing the unresolved native library.


osgiCfgInit script

The Equinox OSGi framework is used to manage class loading and relationships between server component bundles. In some cases, the cached bundle data, which is maintained on a per-profile basis and has a separate cache at the WAS_HOME level for installation-wide processes, can become out of sync with the actual binaries on the server. You can use the osgiCfgInit script to clear and recreate the OSGi cache.

You should run the osgiCfgInit script on the command line from the WAS_HOME/bin or *user_install_root/bin* directory. The behavior of the script depends on the directory from which you run the script. If you run the script from a profile-level bin directory, the script clears the OSGi cache for all servers within that profile. If you run the script from the WAS_HOME/bin directory, the script clears the OSGi cache for all servers within the default profile.

Note: Before you run the osgiCfgInit script, stop the server on which the script will be run. If you run this script on a server that is active, the server might have problems trying to read or update the cache after the script is finished.

Syntax

The syntax for this script is as follows: 

```
osgiCfgInit.sh [options]
```

Options

The following options are available for the osgiCfgInit script:

-all


The script clears the caches of all servers in the installation, as well as the WAS_HOME cache.

-washome

The script clears the cache at the WAS_HOME level.

Usage scenario

To clear the cache, execute the command as follows:



```
./osgiCfgInit.sh [-all|-washome]
```

If the script completes successfully, the message "OSGi cache successfully cleaned for *location*." displays on the command line.

Class loader viewer service settings

Use this page to configure the server to start the class loader viewer service when the server starts. The Class Loader Viewer helps you diagnose problems with class loaders.

To view this administrative console page, click **Servers > Server Types > WebSphere application servers > *server_name* > Class loader viewer service**.

Class loaders find and load class files. For a deployed application to run properly, the class loaders that affect the application and its modules must be configured so that the application can find the files and

resources that it needs. Diagnosing problems with class loaders can be complicated and time-consuming. To diagnose and fix the problems more quickly, enable the class loader viewer service on this page and then use the console Class loader viewer to examine class loaders and the classes loaded by each class loader. Click **Troubleshooting** > **Class loader viewer** to access the Class loader viewer in the console.

Enable service at server startup

Specifies whether or not the server attempts to start the class loader viewer service when the server starts.

The default is not to start the class loader viewer service.

Enterprise application topology

Use this page to see where modules reside in a topology of enterprise applications. Knowing where a module resides helps you to determine which class loader loaded a module and to diagnose problems with class loaders.

To view this administrative console page, click **Troubleshooting** > **Class loader viewer**. This page lists all installed applications and their modules in a tree view. The modules can be web modules (.war files) or enterprise bean (EJB) modules (.jar files).

When deploying an application to a server or starting an application, you might encounter problems related to class loaders. Use the console pages accessed from this page to troubleshoot errors such as the following:

- ClassCastException
- ClassNotFoundException
- NoClassDefFoundException
- UnsatisfiedLinkError

You can use the Class loader viewer console pages without having to restart or manipulate the application.

Enterprise applications topology

Displays a tree hierarchy of applications installed on a server and lists the module files in the class paths of the applications.

Expand the hierarchy for an application to see what web modules (.war files) and EJB modules (.jar files) are in the application class path.

Click on a module name to examine the class loaders of the module.

Class loader viewer settings

Use this page to examine the class loaders visible to a web module (.war file) or enterprise bean (.ejb file) in an installed enterprise application. This page helps you to determine which class loaders loaded files of a module and to diagnose problems with class loaders.

To view this administrative console page, click **Troubleshooting** > **Class loader viewer** > *module_name*.

To learn more about classes used by the module and their class loaders, click a button:

Table 33. Class loader viewer buttons. Click a button to access information about classes.

Button	Resulting action
Export	Opens a dialog that enables you to view or save the class loader information on this page in an XML file.

Table 33. Class loader viewer buttons (continued). Click a button to access information about classes.

Button	Resulting action
Table View	<p>Displays the Table view page, which provides information about class loaders visible to the module in an HTML table format for each class loader. Such information includes:</p> <p>Delegation Whether the class loader delegates a load operation to its immediate parent before searching its local classpath for a class or resource</p> <p>Classpath The local classpath, which includes the paths over which the class loader searches for classes and resources, excluding the classpaths of any parent class loaders.</p> <p>Classes The names of classes loaded by the class loader</p>
Search	<p>Displays the Search page, on which you can search class loaders for the following:</p> <ul style="list-style-type: none"> • Specific strings • Specific .jar files • The names of files in a specific directory • The names of files loaded by a specific class loader

Class Loader

Displays a hierarchy of class loaders that affect the loading of classes used by the web or EJB module. The **Hierarchy** tab displays the class loaders in a tree hierarchy. The **Search Order** tabs lists the class loaders in the order in which the runtime environment uses them to find and load classes.

Expand a hierarchy of class loaders to view the following:

- Class loader names
- Arrows that point upwards beside class loader names, indicating that requests can go to a parent class loader only and not go to a child class loader
- The names of classes that are loaded by a class loader
- The paths of property files and .jar files used by the classes

The following class loaders might be in a hierarchy:

Table 34. Class loader name descriptions. Class loaders that might be in the hierarchy of class loaders.

Class loader name	Description
JDK Extension Loader	The JDK extensions class loader is a composite class loader that is comprised of the Java virtual machine (JVM) bootstrap class loader, the JVM extensions class loader and the JVM system class loader, which load the core SDK classes and resources as well as classes and resources visible on the JVM classpath.
WAS Extension Class Loader	The WAS Extension Class Loader loads the WebSphere Application Server classes, stand-alone resource classes, custom service classes, and custom registry classes. At bootstrap, this class loader uses the <code>ws.ext.dirs</code> system property to determine the path that is used to load classes. Each directory in the <code>ws.ext.dirs</code> class path and every .jar file or compressed .zip file in these directories is added to the class path used by this class loader.
WAS Compound Class Loader	The WAS Compound Class Loaders load classes and resources of enterprise archive (EAR) modules, web application archive (WAR) modules, and server-associated shared libraries. Under default class loader policies, an instance of a WAS Compound Class Loader exists for each running EAR and WAR module and for each class loader defined in the server configuration.

Click on **Classes** to view a list of classes loaded by a class loader.

The class loader viewer service must be enabled to view the list of classes.

Search settings

Use this page to search for information about class loaders visible to a web module (.war file) or enterprise bean (.ejb file) in an installed enterprise application. This page helps you diagnose problems with class loaders.

To view this administrative console page, click **Troubleshooting** > **Class loader viewer** > *module_name* > **Search**.

On the Search page, you can search class loaders for the following:

- Specific strings
- Specific .jar files
- The names of files in a specific directory
- The names of files loaded by a specific class loader

Search type

Specifies the type of items in which to search for the string.

*Table 35. Search type fields. Type searchable information in a field and click **Go**.*

Search type	Instructions and resulting action
Class/Package	In the Search terms field, type a class name or package name. After you select this search type and click Go , the program searches class loaders for a class or package name. The program displays a list of classes and packages that have the string in their name.
JAR/Directory	In the Search terms field, type a .jar file name or directory name. After you select this search type and click Go , the program searches class loaders for a .jar file or directory name. The program displays a list of .jar files that have the string in their name and of all files in directories that have the string in their name.

Search terms

Specifies the string to be found in the items searched.

The search is case-sensitive. If the search string is `classname`, the string `ClassName` is not found.

The search matches the entire string. If the search type is **JAR/Directory** and the search string is `C:/WebSphere/AppServerd0603.185/java/jre/lib/ext/CmpCrmf.jar`, the entire path of the JAR file is matched. If the search type is **JAR/Directory** and the search string is `Cmp`, the string `Cmp` is not found.

The search supports limited regular expressions. It supports the wildcard characters asterisk (*), question mark (?), and percent sign (%). The wildcard characters * and % match zero or more characters; ? matches exactly one character.

Table 36. Search strings with wildcard characters. Use the example search strings to see what items result from searches.

Search string	Resulting matches
Cmp	Items that have Cmp in their name
Cmp.jar	Items that have Cmp in their name and that end in .jar
%Cmp%	Items that have Cmp in their name
%Cmp%.jar	Items that have Cmp in their name and that end in .jar

Table 36. Search strings with wildcard characters (continued). Use the example search strings to see what items result from searches.

Search string	Resulting matches
*Cmp?rmf.jar	Items that have a name with any characters before Cmp, then any one character, and then rmf.jar

The search supports full regular expressions if the value for the search string starts and ends with a forward slash (/).

Table 37. Search strings with regular expressions. Use the example search strings to see what items result from searches.

Search string	Resulting matches
/. *Cmp.*/	Items that contain any character before and after Cmp in their name
/. *Cmp.*\.jar/	Items that have Cmp in their name and that end in .jar
/. *Cmp?rmf\.jar/	Items that have a name with any characters before Cmp, then any one character, and then rmf.jar
/. *\d\.jar/	Items with a name that ends in a number followed by .jar

Chapter 13. Configuring the hang detection policy

The hang detection option for WebSphere Application Server is turned on by default. You can configure a hang detection policy to accommodate your applications and environment so that potential hangs can be reported, providing earlier detection of failing servers. When a hung thread is detected, WebSphere Application Server notifies you so that you can troubleshoot the problem.

Before you begin

A common error in Java Platform, Enterprise Edition (Java EE) applications is a hung thread. A hung thread can result from a simple software defect (such as an infinite loop) or a more complex cause (for example, a resource deadlock). System resources, such as CPU time, might be consumed by this hung transaction when threads run unbounded code paths, such as when the code is running in an infinite loop. Alternately, a system can become unresponsive even though all resources are idle, as in a deadlock scenario. Unless an end user or a monitoring tool reports the problem, the system may remain in this degraded state indefinitely.

Using the hang detection policy, you can specify a time that is too long for a unit of work to complete. The thread monitor checks all managed threads in the system (for example, web container threads and object request broker (ORB) threads) . Unmanaged threads, which are threads created by applications, are not monitored. For more information read “Hung threads in Java Platform, Enterprise Edition applications” on page 175.

About this task

The thread hang detection option is enabled by default. To adjust the hang detection policy values, or to disable hang detection completely:

Procedure

1. From the administrative console, click **Servers > Application Servers > server_name**
2. Under Server Infrastructure, click **Administration > Custom Properties**
3. Click **New**.
4. Add the following properties:

Information	Description
Name	com.ibm.websphere.threadmonitor.interval
Value	The frequency, in seconds, at which managed threads in the selected application server will be interrogated.
Default	180 seconds (three minutes)

Information	Description
Name	com.ibm.websphere.threadmonitor.threshold
Value	The length of time, in seconds, in which a thread can be active before it is considered hung. Any thread that is detected as active for longer than this length of time is reported as hung.
Default	600 seconds (ten minutes)

Information	Description
Name	com.ibm.websphere.threadmonitor.false.alarm.threshold

Information	Description
Value	The number of times (T) that false alarms can occur before automatically increasing the threshold. It is possible that a thread that is reported as hung eventually completes its work, resulting in a false alarm. A large number of these events indicates that the threshold value is too small. The hang detection facility can automatically respond to this situation: For every T false alarms, the threshold T is increased by a factor of 1.5. Set the value to zero (or less) to disable the automatic adjustment.
Default	100

Information	Description
Name	com.ibm.websphere.threadmonitor.dump.java
Value	<p>IBM i Set to true to cause a javacore to be created when a hung thread is detected and a WSVR0605W message is printed. The threads section of the javacore can be analyzed to determine what the reported thread and other related threads are doing. See the topic <i>Dumping threads in server processes using scripting</i> for a description of the content of this dump.</p> <p>This property calls the dumpThreads function. Set to true to cause a javacore, a heapdump, and a TDUMP to be created when a hung thread is detected and a WSVR0605W message is printed. The threads section of the javacore can be analyzed to determine what the reported thread and other related threads are doing. See the topic <i>Dumping threads in server processes using scripting</i> for a description of the contents of these dumps, and how to disable the creation of one or more of these dumps.</p> <p>Set to an integer value in the range 1 through Integer.MAX_VALUE to cause a javacore to be created when a hung thread is detected and a WSVR0605W message is printed. The integer value indicates the maximum number of javacores that will be created.</p>
Default	false (0)

Information	Description
Name	com.ibm.websphere.threadmonitor.dump.stack
Value	Set to true to cause a stack trace to be printed when a hung thread is detected and a WSVR0605W message is printed.
Default	true

To disable the hang detection option, set the `com.ibm.websphere.threadmonitor.interval` property to less than or equal to zero.

- Optional: If you want to monitor the activity of threads on which system alarms execute, add the following JVM generic arguments to the server settings.

Information	Description
Name	-Dcom.ibm.websphere.alarmthreadmonitor.generate.javacore
Value	Set to any value to cause a javacore dump to be created when an hung system alarm thread is detected. The threads section of the javacore dump can be analyzed to determine what the reported thread and other related threads are doing.
Default	Unset

Information	Description
Name	com.ibm.websphere.alarmthreadmonitor.checkinterval.millis
Value	The frequency, in milliseconds, at which system alarm threads are interrogated. Set the value to zero to disable system alarm hung thread detection. The maximum interval is 600000 (10 minutes).
Default	10000 (10 seconds)

Information	Description
Name	-Dcom.ibm.websphere.alarmthreadmonitor.threshold.millis

Information	Description
Value	Set to any value integer between 10000 and 600000 (10 minutes). This argument is used to specify the length of time, in milliseconds, that a system alarm thread can be active before it is considered non-responsive. Any system alarm thread that is detected as inactive for longer than this length of time is reported as hung.
Default	10000 (10 seconds)

To add these arguments to the server settings, complete the following actions:

- a. Under Server Infrastructure on the server settings page in the administrative console, click **Java and process management > Process definition**.
 - b. Select **Java virtual machine**.
 - c. Add the arguments to the JVM generic arguments section.
6. Click **Apply**.
 7. Click **OK**.
 8. Save the changes.
 9. Restart the Application Server for the changes to take effect.

Hung threads in Java Platform, Enterprise Edition applications

WebSphere Application Server monitors thread activity and performs diagnostic actions if one has become inactive.

When WebSphere detects that a thread has been active longer than the time defined by the thread monitor threshold, the application server takes the following actions:

- Logs a warning in the WebSphere Application Server log that indicates the name of the thread that is hung and how long it has already been active. The following message is written to the log:

```
WSVR0605W: Thread threadname has been active for
hangtime and may be hung. There are totalthreads
threads in total in the server that may be hung.
```

where: *threadname* is the name that appears in a JVM thread dump, *hangtime* gives an approximation of how long the thread has been active and *totalthreads* gives an overall assessment of the system threads.

- Issues a Java Management Extensions (JMX) notification. This notification enables third-party tools to catch the event and take appropriate action, such as triggering a JVM thread dump of the server, or issuing an electronic page or email. The following JMX notification events are defined in the `com.ibm.websphere.management.NotificationConstants` class:
 - `TYPE_THREAD_MONITOR_THREAD_HUNG` This event is triggered by the detection of a (potentially) hung thread.
 - `TYPE_THREAD_MONITOR_THREAD_CLEAR` This event is triggered if a thread that was previously reported as hung completes its work. Consult the section on false alarms for more information.
- Triggers changes in the performance monitoring infrastructure (PMI) data counters. These PMI data counters are used by various tools, such as the Tivoli Performance Viewer, to provide a performance analysis.
- Triggers changes in the performance monitoring infrastructure (PMI) data counters. These PMI data counters are used by various tools, such as the Tivoli Performance Viewer, to provide a performance analysis.

For additional information about performance monitoring and Tivoli Performance Viewer, see the chapter Monitoring performance with Tivoli Performance Viewer in the *Tuning guide* PDF book

False Alarms

If the work actually completes, a second set of messages, notifications and PMI events is produced to identify the false alarm. The following message is written to the log:

```
WSVR0606W: Thread threadname was previously reported to be
hung but has completed. It was active for approximately hangtime.
There are totalthreads threads in total in the server that still
may be hung.
```

where *threadname* is the name that appears in a JVM thread dump, *hangtime* gives an approximation of how long the thread has been active and *totalthreads* gives an overall assessment of the system threads.

Automatic adjustment of the hang time threshold

If the thread monitor determines that too many false alarms are issued (determined by the number of pairs of hang and clear messages), it can automatically adjust the threshold. When this adjustment occurs, the following message is written to the log:

```
WSVR0607W: Too many thread hangs have been falsely reported. The hang
threshold is now being set to thresholdtime.
```

where: *thresholdtime* is the time (in seconds) in which a thread can be active before it is considered hung.

You can prevent WebSphere Application Server from automatically adjusting the hang time threshold. See Chapter 13, “Configuring the hang detection policy,” on page 173

System Alarms

An application server monitors the activity of threads on which system alarms execute. When a system alarm thread has been active longer than the time defined by the alarm thread monitor threshold, the application server logs the following warning in the system log. This message indicates the name of the thread that is not responding, the length of time that the thread has already been active, and the exception stack of the thread, which identifies the system component.

```
UTLS0008W: The alarm thread threadname has been active for n
milliseconds and may be hung. totalthreadsthreadstack
```

In this message, *threadname* is the name that appears in a JVM thread dump, *n* is approximately how long the thread was active, *totalthreads* is an overall assessment of the system threads, and *threadstack* is the exception stack of the thread.

If the alarm work eventually completes, the following message is written to the system log. This message indicates thread that produced the false alarm.

```
UTLS0009W: Alarm Thread threadname was previously reported to be hung but has
completed. It was active for approximately n milliseconds.
```

In this message, *threadname* is the name that appears in a JVM thread dump, and *n* is approximately how long the thread was active.

Typically, system alarms do not process heavy loads because such activity might slow the processing of later system alarms, which in turn might impact server behavior. The UTLS0008W message is intended to help IBM Support personnel investigate problems potentially caused by system alarm behavior.

All of the system alarms share a common alarm thread pool. The properties which govern the monitoring of this thread pool can be tuned using the administrative console. You can reduce the frequency at which WebSphere generates alarm hung thread messages by adjusting the alarm thread monitor check interval or threshold. See the topic *Configuring the hang detection policy* for a description of how to change these settings.

Example: Adjusting the thread monitor to affect server hang detection

The hang detection policy affects how the application server responds to a thread that is not being processed correctly.

You can adjust the thread monitor settings by using the wsadmin scripting interface. These changes take effect immediately, but do not persist to the server configuration, and are lost when the server is restarted. The following script provides an example of how to adjust the properties for the thread monitor using the wsadmin tool:

```
# Read in the interval, threshold, false alarm from the command line
set interval [lindex $argv 0]
set threshold [lindex $argv 1]
set adjustment [lindex $argv 2]

# Get the object name of the server you want to change the values on
set server [$AdminControl completeObjectName "type=Server,*"]

# Read in the interval and print to the console
set i [$AdminControl getAttribute $server threadMonitorInterval]

# Read in the threshold and print to the console
set t [$AdminControl getAttribute $server threadMonitorThreshold]

# Read in the false alarm adjustment threshold and print to the console
set a [$AdminControl getAttribute $server threadMonitorAdjustmentThreshold]

# Set the new values using the command line parameters
$AdminControl setAttribute $server threadMonitorInterval ${interval}

$AdminControl setAttribute $server threadMonitorThreshold ${threshold}

$AdminControl setAttribute $server threadMonitorAdjustmentThreshold ${adjustment}
```

Chapter 14. Working with troubleshooting tools

WebSphere Application Server includes a number of troubleshooting tools that are designed to help you isolate the source of problems. Many of these tools are designed to generate information to be used by IBM Support, and their output might not be understandable by the customer.

About this task

This section only discusses tools that are bundled with the WebSphere Application Server product. A wide range of tools which address a variety of problems is available from the WebSphere Application Server Technical Support website.

Procedure

1. Select the appropriate tool for the task. For more information on the capacities of the supplied troubleshooting tools, see the relevant articles in this section.
2. Run the tool as described in the relevant article.
3. Contact IBM Support for assistance in deciphering the output of the tool. 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. For the last minute updates, limitations, and known problems, see the release notes. Before opening a PMR, see the Must gather page.
4. Use the IBM Support Assistant to help find and use various IBM Support resources, such as updated documentation and problem determination tools.

Gathering information with the collector tool (deprecated)

The collector tool gathers information about your WebSphere Application Server installation and packages it in a Java archive (JAR) file that you can send to IBM Customer Support to assist in determining and analyzing your problem. Information in the JAR file includes logs, property files, configuration files, operating system and Java data, and the presence and level of each software prerequisite.

Before you begin

The sort of information that you gather is not something that most people use. In fact, the collector tool packages its output into a JAR file. IBM includes the collector tool in the product code, along with other tools that help capture the information that you must provide when reporting a problem. The collector tool is part of a strategy of making problem reporting as easy and complete as possible.

There are two phases of using the collector tool. The first phase runs the collector tool on your WebSphere Application Server product and produces a Java archive (JAR) file. The IBM Support team performs the second phase, which is analyzing the Java archive (JAR) file that the collector program produces. The collector program runs to completion as it creates the JAR file, despite any errors that it might find like missing files or invalid commands. The collector tool collects as much data in the JAR file as possible.

The collector tool is a Java application that requires a Java SE Runtime Environment 6 (JRE6) to run.

About this task

The tool is within the installation root directory for WebSphere Application Server, Express. But you run the tool from a working directory that you create outside of the installation root directory. This procedure describes both of those steps and all of the other steps for using the tool and reporting the results from running the tool.

There are two ways to run the collector tool. Run the collector tool to collect summary data or to traverse the system to gather relevant files and command results. The collector tool produces a Java archive (JAR) file of information needed to determine and solve a problem. The collector summary option produces a lightweight collection of version and other information that is useful when first reporting the problem to IBM Support. Run the collector tool from the root user or from the administrator user to access system files that contain information about kernel settings, installed packages, and other vital data.

The tool collects information about the default profile if you do not use the optional parameter to identify another profile.

Procedure

IBM i Run the collector tool.

1. Log on to the system with a user profile that has all object (*ALLOBJ) special authority.
2. Make a working directory where you can start the collector program.
3. Run the **STRQSH** command from the CL command line to prepare to run the collector program.
4. Make the working directory the current directory.
5. Run the following command from Qshell:

```
cd workingDirectory
```

The collector program writes its output JAR file to the current directory. The program also creates and deletes a number of temporary files in the current directory. Creating a work directory to run the collector program avoids naming collisions and makes cleanup easier. You cannot run the collector tool in a directory under the installation root directory for WebSphere Application Server, Express.

6. Run the collector program by entering the fully qualified command from the command line of the working directory.

- Run the following command from Qshell:

```
app_server_root/bin/collector
```

- Use the following command to gather data from a specific profile that might not be the default profile.

- Run the following command from Qshell:

```
app_server_root/bin/collector -profileName profile_name
```

7. Optional: You can also run the collector tool from the profile's root directory instead of the *app_server_root/bin/* directory.

Run the following command from Qshell:

```
profile_root/bin/collector
```

You should get the same output if you run the collector tool from the bin directory of *profile_root* as you would running it from *app_server_root*.

Issuing the command from the profile also runs the setupCmdLine file in the profile's bin directory. This file sets an environment parameter that the collector uses to determine which profile's data to collect.

Results

The collector program creates the Collector.log log file and an output JAR file in the current directory.

The name of the JAR file is composed of the host name, cell name, node name, and profile name:

```
host_name-cell_name-node_name-profile_name.JAR
```

The Collector.log log file is one of the files collected in the *host_name-cell_name-node_name-profile_name.JAR* file.

What to do next

Send the *host_name-cell_name-node_name-profile_name*.JAR file to IBM Support for analysis.

Collector tool output

Use the collector tool to gather and analyze output from WebSphere Application Server.

The first step in using the collector tool on your WebSphere Application Server product is to run the tool to produce a Java archive (JAR) file as output. The second step in using the collector tool is to analyze its output. The preferred method of performing this analysis is to send the JAR file to IBM Support for analysis. However, you can use this topic to understand the content of the JAR file if you perform your own analysis.

You can view the files contained in the JAR file without extracting the files from the JAR file. However, it is easier to extract all files and view the contents of each file individually. To extract the files, use one of the following commands:

- `jar -xvf WASenv.jar`
- `unzip WASenv.jar`

Wasenv.jar stands for the name of the JAR file that the collector tool creates.

The JAR file contains:

- A collector tool log file, `collector.log`
- Copies of stored WebSphere Application Server files and their full paths that are located under directory root in the JAR file
- Java information in a directory named Java
- A JAR file manifest

Tips and suggestions

- Unzip the JAR file to an empty directory for easy access to the gathered files and for simplified cleanup.
- Check the `collector.log` file for errors:
 - Some errors might be normal or expected. For example, when the collector attempts to gather files or directories that do not exist for your specific installation, it logs an error about the missing files.
 - A non-zero return code means that a command that the collector tool attempted to run does not exist. This might be expected in some cases. If this type of error occurs repeatedly, there might actually be a problem.

collector command - summary option

WebSphere Application Server products include an enhancement to the collector tool beginning with Version 5.0.2, known as the *collector summary option*.

The collector summary option helps you communicate with WebSphere Application Server technical staff at IBM Support. Run the collector tool with the `-Summary` option to produce a lightweight text file and console version of some of the information in the Java archive (JAR) file that the tool produces without the `-Summary` parameter. You can use the collector summary option to retrieve basic configuration and prerequisite software level information when starting a conversation with IBM Support.

The collector summary option produces version information for the WebSphere Application Server product and the operating system as well as other information. It stores the information in the `Collector_Summary.txt` file and writes it to the console. You can use the information to answer initial questions from IBM Support or you can send the `Collector_Summary.txt` file directly to IBM Support.

IBM i When running the collector tool on WebSphere Application Server for IBM i, an additional file named `WAS_Collection_timestamp.html` is created in the `logs` directory of the profile where the collector tool runs. If you do not specify the `-profileName` option, the profile is the default profile. The HTML file is sent to IBM Support when you initiate a PMR.

Run the **collector** command to create the JAR file if IBM Support needs more information to solve your problem.

To run the collector summary option, start from a temporary directory outside of the WebSphere Application Server product installation root directory and enter one of the following commands:

- **IBM i** `app_server_root/bin/collector -Summary [-profileName profileName]`

The `-profileName` option specifies the profile to summarize.

First failure data capture (FFDC)

When a failure occurs during product runtime processing, the (FFDC) feature instantly collects information about the events and errors that lead up to the failure. The captured data can then be used to analyze the problem. After a maximum number of days, these files are automatically deleted from your system.

IBM i After the information is collected, and saved in a log file, FFDC returns control to the affected engines.

By default, a FFDC log file is automatically purged seven days it is created. You can configure the amount of days between purges if you are concerned about the amount of space that the FFDC log files are using.

Two FFDC implementations are provided in the product:

- The WebSphere FFDC, which is the legacy FFDC implementation. This FFDC can only be used in WebSphere products.
- The IBM FFDC, which is a more componentized, and more generic implementation that depends solely on the JDK. This FFDC can be used in client processes and by non-WebSphere products, because it is pluggable with non-WebSphere data collectors, formatters, providers, and listeners.

Both of these FFDC implementations support the `OnDirProvider` type functionality that is configurable using the `com.ibm.ffdc.log` Java environment variable. The `OnDirProvider` functionality includes a built-in provider that stores incidents as separate files in a directory, along with a separate summary file. The `com.ibm.ffdc.log` Java environment variable can be set to the following values:

- **<file_name>**, where *file_name* is either be the name of a single file or a directory path.
 - If *file_name* exists and is the name of a single file, all of the incident and summary reporting information that FFDC collects is appended into that file.
 - If *file_name* exists and is a directory path, whenever an incident occurs, a new file is created in that directory and all of the incident and summary reporting information for that incident is written into this newly created file. The incident is also added to the summary report in this directory.
 - If *file_name* ends in a file separator (`\` or `/`). but a file with the specified name does not exist, a directory called *file_name* is created. Then, whenever an incident occurs, a new file is created in that directory and all of the incident and summary reporting information for that incident is written into this newly created file. The incident is also added to the summary report in this directory.
 - If *file_name* does not end in a file separator (`\` or `/`). and a file with the specified name does not exist, a single file is created and given the specified name. All of the incident and summary reporting information that FFDC collects is then appended into that file.
- `System.out`, which appends the incidents and summary report information to the `stdout` output stream.
- `System.err`, which append the incidents and summary report information to the `stderr` output stream. `System.err` is the default value for the `com.ibm.ffdc.log` Java environment variable.
- `Suppress`, which causes all FFDC collected information to be discarded.

Specifying a value for the `com.ibm.ffdc.log` Java environment variable is the only configuration change that you need to make to exploit the `OnDirProvider` functionality. The new FFDC also provides mechanisms to overly choose this provider or to use your own WebSphere provider.

Starting in WebSphere Application Server V8, this same variable can be used to redirect the FFDC incidents and summary created in the product. The only acceptable value for that, however, is a directory. So it is recommended that, if this variable is used in a WebSphere server environment, that the specification end in a File separator (`/` or `\\`) to avoid conflicts.

gotcha: If the default setting for automatic purging of FFDC information is too long for your environment, see the topic `Configuring first failure data capture log file purges` for a description of how you can modify the length of time that the FFDC information is retained on your system.

Configuring first failure data capture log file purges

The first failure data capture (FFDC) log file saves information that is generated from a processing failure. These files are deleted after a maximum number of days has passed. The captured data is saved in a log file for analyzing the problem.

Before you begin

IBM i The first failure data capture (FFDC) feature preserves the information that is generated from a processing failure and returns control to the affected engines. The captured data is saved in a log file for analyzing the problem. FFDC is intended primarily for use by IBM Support. FFDC instantly collects events and errors that occur during the product run time. The information is captured as it occurs and is written to a log file that can be analyzed by IBM Support personnel.

The FFDC configuration properties files are located in the properties directory under the Application Server product installation. You must set the `ExceptionFileMaximumAge` property to the same value in all three files: `ffdcRun.properties`, `ffdcStart.properties`, and `ffdcStop.properties`. You can set the `ExceptionFileMaximumAge` property to configure the days between purging the FFDC log files. The value of the `ExceptionFileMaximumAge` property must be a positive number. The FFDC feature does not affect the performance of the Application Server product.

About this task

Perform the following steps to configure the number of days between the FFDC log file purges. The value is in days.

Procedure

1. Open the `ffdcRun.properties` file.

IBM i The file is located in the `profile_root/properties` directory.
2. Change the value for the `ExceptionFileMaximumAge` property to the number of days between the FFDC log file purges. The value of the `ExceptionFileMaximumAge` property must be a positive number. The default is seven days. For example, `ExceptionFileMaximumAge = 3` sets the default time to three days. The FFDC log file is purged after three days.
3. Save the `ffdcRun.properties` file and exit.
4. Repeat the previous steps to modify the `ffdcStart.properties` and `ffdcStop.properties` files.

Results

IBM i The FFDC file management function removes the FFDC log files that have reached the maximum age and generates a message in the `SystemOut.log` file.

Note: This topic references one or more of the application server log files. As a recommended alternative, you can configure the server to use the High Performance Extensible Logging (HPEL) log and trace infrastructure instead of using SystemOut.log , SystemErr.log, trace.log, and activity.log files on distributed and IBM i systems. You can also use HPEL in conjunction with your native z/OS logging facilities. If you are using HPEL, you can access all of your log and trace information using the LogViewer command-line tool from your server profile bin directory. See the information about using HPEL to troubleshoot applications for more information on using HPEL.

Chapter 15. Working with Diagnostic Providers

Diagnostic Providers enable you to query the startup configuration, current configuration, and current state of a diagnostic domain. In addition, Diagnostic Providers can also provide access to any self diagnostic tests that are available from a diagnostic domain.

About this task

The Diagnostic Provider Utility is a simple front end in the administration console that presents the available set of Diagnostic Providers and enables you to work with them.

Procedure

Learn about Diagnostic Providers

Diagnostic Providers

Diagnostic Providers are a quick method for viewing configuration and the current state of individual components within an application server environment.

WebSphere Application Server components can be considered as being divisible into *diagnostic domains*. A diagnostic domain refers to a set of classes within the component that share a set of diagnostics. Some larger components might have multiple diagnostic domains. For example, the Connection Manager logically consists of multiple data sources and connection factories that each have separate diagnostic domains.




Figure 7. Diagnostic providers. This image shows the relationships between the parts that make up the Diagnostic Provider (DP) utility.

A single diagnostic domain receives its diagnostic services from a Diagnostic Provider MBean. The Diagnostic Provider MBean enables you to query the startup configuration, current configuration, and current state of the diagnostic domain. In addition, Diagnostic Provider MBeans can also provide access to any self diagnostic tests that are available from the diagnostic domain. Some characteristics of Diagnostic Provider MBeans include:

- Diagnostic Provider MBeans are Java Management Extensions (JMX) MBeans
- Diagnostic Provider MBeans all implement a DiagnosticProvider interface which includes methods for configuration dumps, state dumps, and self diagnostic tests
- Diagnostic Provider MBeans provide a way to expose information about running components so administrators can more easily debug problems related to those components. As with other MBeans running in WebSphere Application Server, they can be accessed from JMX client code, or through the *wsadmin* tool.

Diagnostic Provider MBeans are efficient at delivering Java object representations of configuration, state, and self test information. This is good for when programs interact. For human users to access the information, WebSphere Application Server provides a set of facilities to extend the value of Diagnostic Provider MBeans.

The Diagnostic Service MBean

provides methods to convert Diagnostic Provider MBean output into human readable formats. The Diagnostic Service MBean also provides some methods to facilitate looking up the Diagnostic Provider MBeans on the same server as the Diagnostic Service MBean. For administrators that want to access diagnostic data from a command line, the *wsadmin* tool can be used directly with the Diagnostic Service MBean to get formatted results

The Diagnostic Provider utility

a set of panels included in the WebSphere Application Server administration console through which administrators can interact with Diagnostic Provider MBeans. The Diagnostic Provider utility is a simple front end in the administration console that presents the available set of Diagnostic Provider MBeans present on each managed server, and provides a means to execute and view the results of configuration dumps, state dumps, and diagnostic self tests.

The purpose of Diagnostic Providers

Diagnostic Providers give you more information for quickly discovering and diagnosing system problems. The following scenario contrasts the experience of an administrator working with a component that does not have a Diagnostic Provider to one that does.

When the administrator works with a component that is without a Diagnostic Provider, the events are as follows:

1. A log entry indicates that a particular component is experiencing a problem.
2. The system administrator sees the log entry through the runtime messages panel.
3. The system administrator cannot tell what is wrong, so calls IBM support for assistance, with a potentially ill-defined problem.

When the administrator works with a component with a Diagnostic Provider, and the Diagnostic Provider ID is registered with the component's logger, the situation changes as follows:

1. A log entry that contains a Diagnostic Provider ID (DPID) indicates that something has gone wrong in a specific component.
2. The system administrator sees the log entry through the runtime messages panel.
3. The administrator clicks a button on the runtime message panel to execute a state dump or a configuration dump, or to be taken to the list of component self tests.
4. From the self test, the administrator is warned that the component is configured in a way that could lead to poor performance or failures.

Furthermore, when the administrator works with a component with a Diagnostic Provider, and the Diagnostic Provider ID is **not** registered with the component's logger, the situation might unfold like this:


1. A log entry which doesn't contain a DPID indicates that something has gone wrong in a component.
2. The system administrator sees the log entry through the runtime messages panel.
3. The system administrator uses the administrative console to navigate through the available set of Diagnostic Providers and selects one that sounds appropriate.
4. He runs a configuration dump, a state dump, or a self diagnostic test against the Diagnostic Provider to collect information about the component.
5. From the state dump, the administrator is able to notice that the component state is not what would be expected for its workload.
6. The administrator works with the test team to determine which of the flows is causing the state of the component to diverge from what is expected (as evidenced by repeated execution of the state dump).

Diagnostic Provider IDs

A Diagnostic Provider ID (DPID) is the unique address of a Diagnostic Provider MBean. Components that have associated Diagnostic Provider MBeans can include the DPID in their log entries.

Diagnostic Provider IDs are implemented in WebSphere Application Server as Java Management Extensions (JMX) *MBean ObjectNames*, and can be used at JMX MBean servers to look up Diagnostic Provider MBeans.

By including the String representation of the DPID in each logged message, the message can be tracked back to the Diagnostic Provider related to the component. A method is provided to associate Diagnostic Provider IDs with Loggers (from the *java.util.logging* logging API).



The previous diagram shows how the use of DPIDs in log entries enables callbacks to the component that originally created the log entry.

1. Shows the component logging with a DPID included in the log entry.
2. The administrator examines the log entry through the Runtime Messages Utility and notices that the entry has a link to a Diagnostic Provider.
3. The administrator uses the link to gain access to the relevant MBean in the Diagnostic Provider Utility .
4. The Diagnostic Provider Utility contacts the Diagnostic Provider MBean to ask for more information.
5. The request for more information is sent back to the source of the original log entry.
6. The response from the Diagnostic Provider is provided in the administration console.

Diagnostic Provider configuration dumps, state dumps, and self tests

The Diagnostic Provider (DP) infrastructure allows for a software component or stack product in the WebSphere Application Server space to expose key information about its configuration, current state, and current ability to perform operations.

The methods that expose this information might be driven as a result of a message put out by the component (by a logger which automatically includes the Diagnostic Provider ID in each message), or might be driven as a result of an overall system health-check when an administrator or automated tool is monitoring the system.

Configuration dumps

A *Configuration dump* is an operation you can perform on a Diagnostic Provider to list the startup or current values of the configuration attributes for the DP. The name for each data item in this dump should reflect its disposition. That is, each item should be called *startup-xxx* or *current-xxx* to show whether this is a startup or current value. The collection of attributes returned from this operation can be thought of as the **payload** of the configuration dump. More information about payloads can be found in “Diagnostic Provider method implementation” on page 194.

You can find several ways to filter the output of a configuration dump in “Diagnostic Provider registered attributes and registered tests” on page 189.

State dumps

A *State dump* is similar to a configuration dump, but it differs in two key areas. First, a state dump displays current information about the operation of a component. An example is a connection pool. A configuration dump can show *DataSource name*, the *minConnections* (configured or current), the *maxConnections*, the *DataBase name*, and so on. A state dump is more likely to show the current connections in use, the high concurrent use count, the number of times the pool has been expanded, the average time between requesting a connection and returning it, and so forth.

State dumps can be impacted by the values in the State Collection Specification. This is a dynamic specification that controls additional data collection that the component can do at runtime. If additional data is being collected, then a State dump might display more information. The same filters and payload information that apply to Configuration dumps (see “Diagnostic Provider registered attributes and registered tests”) apply to State dumps.

Self Diagnostic tests

Self diagnostic tests are non-invasive operations that a Diagnostic Provider exposes. *Non-invasive* means that if they modify anything for the test, the conclusion of the test reverses the modification. These tests give an administrator the option to test simple functions of a component to see if it is able to perform them.

The filters for a self diagnostic test apply to the test itself, not to the output of the test. A typical use of Self Diagnostic tests could be for a pool manager of some sort to pull an object out of the pool and return it to the pool to verify that this operation can still be performed, and with acceptable performance.

Diagnostic Provider registered attributes and registered tests

Each Diagnostic Provider (DP) provides a list of state dump attributes, configuration dump attributes, self tests, and self test attributes. The tests are operations that the DP can perform. The attributes are pieces of information that are available for collection from a Configuration dump, a State dump, or a specific Self Diagnostic test.

Each attribute can be seen as a piece of information with a label on it. Each attribute is also considered to be either *registered* or *not registered*. A registered attribute is one that should be available from one release of WebSphere Application Server to the next. A nonregistered attribute might not be available in its current form in future releases of the product (no commitment has been made).

When performing a Configuration dump, a State dump, or a Self Diagnostic test, an administrator or automatic tooling can request *only registered* values, or *all* values, depending on the needs of the administrator or tool. Note that the option of filtering results is only available through the Diagnostic Provider's Java Management Extension (JMX) MBean interface, which you can access programmatically or through the wsadmin tool.

The DiagnosticProviderRegistration XML file

The DiagnosticProviderRegistration Extensible Markup Language (XML) file is used in conjunction with the method signatures to filter the results of calling the various methods. This XML file defines the configuration information, state information, and self diagnostic tests exposed by the component. In the configuration and state information, the key working unit is referred to as the attribute. Specification of an attribute is as follows:

```
<attribute>
  <id><Regular Expression representing the attribute name></id>
  <descriptionKey><MsgKey into a ResourceBundle for localization of the label></descriptionKey>
  <registered>true</registered>
</attribute>
```

The parts are as follows:

ID: The attribute's name. This name can be expressed with wildcard characters conforming to regular expression syntax. The registered attribute ID is used in the following places:

- Within Diagnostic Provider configuration dump and state dump methods to determine which attributes to return.
- In the administration console to match description keys to attributes returned from a request to a Diagnostic Provider for a configuration dump, state dump, or self diagnostic operation.

As an example, if a configuration dump returns an attribute with ID **cachedServlet-MyServlet-servletPath** to the administration console, the administration console could use the *descriptionKey* corresponding to the attribute registered as `<id>cachedServlet-.*-servletPath</id>` when selecting what description text to put next to this attribute's name and return values.

descriptionKey:

This is a key into a *resourceBundle* for localization.

registered:

This is a boolean qualifying whether this attribute will be available from one release of the software to the next. If registered is **true**, then this attribute should be available in the next release. If registered is **false**, then there is no guarantee that this attribute will continue to exist. Automation should use some caution when handling non-registered attributes.

Specification of a *selfDiagnosticTest* is as follows:

```
<test>
  <id><Regular Expression for the name of the test></id>
  <descriptionKey><MsgKey into a ResourceBundle for localization of the label></descriptionKey>
  <attribute><One or more attributes which will be output from this test></attribute>
</test>
```

The parts are as follows:

ID: Similar to the ID for the attribute, but in this case, describing the test to be performed instead of the attribute to be returned.

descriptionKey:

This is a key into a *resourceBundle* for localization.

Method interfaces

```
public DiagnosticEvent [] configDump(String aAttributeId, boolean aRegisteredOnly);
public DiagnosticEvent [] stateDump(String aAttributeId, boolean aRegisteredOnly);
```

These methods invoke the configuration or state dump on the component, and specify a regular expression filter for the attributes to return as well as filtering the output to include all matching attributes, or only those attributes which are registered. This enables the administrator or automated software driving the method to specify a subset of the overall fields (especially important if many attributes are exposed or if the State Collection Specification increases the amount of data available). The following helper methods are available to assist with filtering the output.

To take a list of Attributes that are available to return, and filter them:

```
public static AttributeInfo [] queryMatchingDPInfoAttributes(String aAttributeId,
  AttributeInfo [] inAttrs, String [] namesToCheck, boolean aRegisteredOnly) {
```

To take a single Attribute that is available to return, and filter it:

```
public static AttributeInfo queryMatchingDPInfoAttributes(String aAttributeId,
  AttributeInfo [] inAttrs, String nameToCheck, boolean aRegisteredOnly) {
```

To go through a populated set of Attribute Information and remove unneeded parts:

```
public static void filterEventPayload(String aAttributeId, HashMap payload) {
```


For details on these messages, please review the API documentation for the **DiagnosticProviderHelper** class. The basic concept is that, once the component knows what attributes are able to be returned, the helper method will determine which of them should be returned based on the regular expression logic and registration boolean.

The selfDiagnostic Method interface here is similar to that of Configdump and Statedump:

```
public DiagnosticEvent[] selfDiagnostic(String aTestId, boolean aRegisteredOnly)
```

The difference is that the first parameter is a regular expression filter for which test to run.

Diagnostic Provider names

In addition to the Diagnostic Provider ID (DPID), each component that implements the Diagnostic Provider interface must have a Diagnostic Provider Name. While the DPID must be unique within the entire WebSphere Application Server domain, the Diagnostic Provider Name need only be unique within the Java Virtual Machine (JVM).

Unlike the Diagnostic Provider ID, which tends to be long and not human-friendly, the Diagnostic Provider Name should be shorter and easier to read. In addition, by convention it should end in *DP*. The Diagnostic Service MBean (see “The simpler interfaces provided by the Diagnostic Service MBean”) can drive methods on a Diagnostic Provider using its name.

The simpler interfaces provided by the Diagnostic Service MBean

All services for a Diagnostic Provider (DP) are also available through a Java Management Extensions (JMX) interface known as the *Diagnostic Service* interface. The Diagnostic Service interface enables administrators to drive methods against DPs using the Diagnostic Provider Name or Diagnostic Provider ID.

When formatted output is requested of the Diagnostic Service, it is localized to the client Locale. This makes the Diagnostic Service MBean ideal for clients using an interface where consuming complex Java objects, such as those returned from the Diagnostic Provider MBeans, is not feasible. An example of such an interface is the wsadmin tool.

The Diagnostic Service interface provides four signatures for each of the key methods available on the Diagnostic Providers (*configDump*, *stateDump*, and *selfDiagnostic*) objects. Because these method signatures look so similar, this example shows it all through the configDump methods. The four Diagnostic Service methods that map to configDump on a Diagnostic Provider are:

```
public DiagnosticEvent [] configDump(String aDPName, String aAttributeIdSpec, boolean aRegisteredOnly)
public DiagnosticEvent [] configDumpById(String aDPid, String aAttributeIdSpec, boolean aRegisteredOnly)
public String [] configDumpFormatted(String aDPName, String aAttributeIdSpec,
    boolean aRegisteredOnly, Locale aLocale)
public String [] configDumpFormattedById(String aDPid, String aAttributeIdSpec,
    boolean aRegisteredOnly, Locale aLocale) {
```

The first two return exactly what the Diagnostic Provider does. The second two methods act as a pass-through to the actual Diagnostic Provider, but they take the array of Diagnostic Events that the Diagnostic Provider returns, and convert it into a more easily consumable *String* array. In addition, these methods handle localizing the output to the appropriate locale. It is important to note that the same method can be driven using the Diagnostic Provider ID or the Diagnostic Provider Name.

Creating a Diagnostic Provider

Use Diagnostic Providers to query the startup configuration, current configuration, and current state of a diagnostic domain. Diagnostic Providers also provide access to any self diagnostic tests that are available from a diagnostic domain.

Before you begin

To complete this task you must have programming knowledge of your system and the proper authorities to perform the following steps.

About this task

The steps that follow outline a general process for creating Diagnostic Providers (DP).

Procedure

1. Determine your diagnostic domain. Look for *configuration* MBeans that control a similar domain in the same component. Extending an existing configuration MBean with a DP interface avoids proliferation of new MBeans and has the benefit that mapping from a diagnostic MBean to a configuration MBean requires no additional information.
2. Determine what configuration attributes you want to expose. Include information that is used to configure your component from the configuration MBeans.
3. Determine what state attributes you want to expose. Anything you might want to know about the state of your component for troubleshooting can go here.
4. Determine what self diagnostic tests you will expose.
5. Determine what test attributes you will return for each self diagnostic.
6. Create your DP registration Extensible Markup Language (XML) file.
7. Create your DP implementation.
 - a. To see an example, refer to “Implementing a Diagnostic Provider” on page 194 and keep in mind that most things that a Diagnostic Provider should do are already done for you in the *DiagnosticProviderHelper* class.
 - b. To ensure that you do not collect unwanted data, add hooks in your component code where you need to collect state data using the *DiagnosticConfig* object.
 - c. Add hooks in your component code where you need to store or be able to access configuration data.
8. Add code to register your DP implementation. Typically, the best place to do this is where your component is initialized.
9. Add Diagnostic Provider IDs (DPID) to your logged messages. Registering a DPID with a logger makes that information available in any messages logged with this logger. This enables fast paths in the DP utility to function on this particular Diagnostic Provider.
 - a. Register your DPID with your loggers (for any of your loggers that you only want to associate a single DPID with).
 - b. When you use multiple DPIDs with the same Logger, you can (instead of registering a single DPID with a Logger) add DPIDs to individual logging calls in the **parm[0]** position. Do not put **{0}** in the corresponding localized messages. It is bad practice to print the DPID in your messages as this would be inconsistent with messages from loggers with statically assigned DPIDs.

Diagnostic Provider Extensible Markup Language

Some conventions to follow for Diagnostic Provider (DP) Extensible Markup Language (XML) declarations.

These guidelines are to help keep your use of Diagnostic Providers (DP) consistent.

- Include the Document type definition (DTD) for your Diagnostic Provider at the top of every DP declaration Extensible Markup Language (XML) file.
- Start all names and name segments with lower case. Use camel case for attribute names. That is, capitalize every initial letter in the name, except the first. For example, *traceCollectionSpec*.
- Indicate hierarchy with dashes. Dashes work better than dots because attribute names are regular expressions. For example, *traceService-traceCollectionSpec*.

- Indicate string dynamic parts to attribute names using an asterisk (*). For example, `vhosts-.*-webgroups-.*-webapps-.*-listeners-filterInvocationListeners`

which would match `vhosts-someHost-webgroups-someGroup-webapps-someApp-listeners-filterInvocationListeners`

- Indicate numeric dynamic parts to attribute names using `[0-9]*`. For example, `vhosts-index-[0-9]*`

which would match `webcontainer-vhosts-index-123`

- If you have a general purpose self diagnostic test that can be run without significant performance cost, name it `general`.

Some tips for configDump implementation

- `configDump` should contain information used to define the component's environment. Some examples are:
 - configuration data set by Java Management Extensions (JMX)
 - configuration from system properties, xml files, and property files
 - configuration information hard-wired and unchanging in code (such as, if a resource adapter implements interface X, or has some static final field Y, then those could indicate aspects of configuration and be included in the `configDump`).
- `configDump` should not contain dynamically registered attributes, such as:
 - a list of registered loggers (this belongs in `stateDump`)
 - a list of servlets in an application (this belongs in `stateDump`).
- `configDump` should be separated into two sections -- *startup* and *current*.
 - All `configDump` attributes must start with either *startup-* or *current-*.
 - The *startup* section details the component's environment at startup time. Startup `configDump` attributes start with *startup-*.
 - The *current* section details the component's environment at the moment the `configDump` is requested. Current `configDump` attributes start with *current-*.

Best practices for configDump

- Group related attributes using an attribute hierarchy (such as, for two attributes about the `traceLog`: `startup-traceLog-rolloverSize=20`, `startup-traceLog-maxNumberOfBackupFiles=1`)
- For information in the current attribute list that refers to the same thing as a startup attribute, the names of both current and startup attributes should match.
- If an attribute has no use following startup, only show it in the startup section (for example, a configuration attribute that contains a file name from which startup data is read).

Choosing a Diagnostic Provider name

To ensure consistency when choosing Diagnostic Provider names to use with your components, you should consider the guidelines that follow.

Diagnostic Provider name guidelines:

- Names must be unique within a Java Virtual Machine (JVM). One Diagnostic Provider name goes uniquely with one Diagnostic Provider ID within a server.
- If necessary, names can contain a dynamic element to help with uniqueness. Of course, the dynamic element should have meaning to the administrator.
- Although not a hard limit, the static part of names should be 16 characters or less.
- The static part of names must follow the class name convention. Start with a capital letter, no spaces, and capitalize each word in the name.

- The static part of names must end with **DP**.
- Valid names contain a static part only, or a static part followed by a dash (-), followed by a dynamic part. Some valid examples:
 - ConnMgrDP-instance_specific_stuff
 - WebContainerDP
 - AdvisorDP
 - NodeAgentDP

Implementing a Diagnostic Provider

To use a Diagnostic Provider you must configure an MBean with the methods and attributes required to handle the data from the application server and client applications.

Before you begin

This task presumes that you have a programming knowledge of the creation of MBeans. For more information about the interaction of MBeans with WebSphere Application Server, refer to topic, *Creating and registering standard, dynamic, and open custom MBeans in the Administering applications and their environment* PDF book.

About this task

The steps that follow outline a general process for implementing a Diagnostic Provider (DP).

Procedure

1. Modify the MBean descriptor Extensible Markup Language (XML). To implement a Diagnostic Provider, you must have an MBean, and the MBean should include this statement in its descriptor XML as a direct child of the MBean element:

```
<parentType type="DiagnosticProvider"/>
```

This defines the operations, attributes, and aggregators necessary for an MBean to be a Diagnostic Provider. If you do not need to have this DP exist in z/OS Controllers, then this XML inclusion handles all z/OS specifics for your MBean.

2. Modify the MBean Implementation. Your MBean should already have a class which instantiates it and registers it with the Java Management Extensions (JMX) server.

The first difference here is that you must define a property in the *Properties* class that is passed to the registration (and becomes part of the *ObjectName*). The property is **diagnosticProvider=true** and it can be added with a line of code such as:

```
MyProps.setProperty(DiagnosticProvider.DIAGNOSTIC_PROVIDER_KEY, DiagnosticProvider.DIAGNOSTIC_PROVIDER_VALUE) ;
```

The second difference is that this class should register this Diagnostic Provider with the Diagnostic Service. A helper method is available to do this:

```
DiagnosticProviderHelper.registerMBeanWithDiagnosticService(DiagnosticProviderPName, DiagnosticProviderId) ;
```

Obviously this must be done after the registration when the *ObjectName* can be retrieved into the *DiagnosticProviderId* string.

3. Implement the Diagnostic Provider methods.

Diagnostic Provider method implementation

To create a Diagnostic Provider (DP) you must have an MBean that includes the required methods in its deployment Extensible Markup Language (XML) file. These methods define the operations, attributes, and aggregators necessary for an MBean to be a Diagnostic Provider.

Adding these methods can be accomplished by adding the *parentType* directive to your existing XML file (see "Implementing a Diagnostic Provider"), or by including the operations directly into your deployment

XML file. The definitions needed are included in “Diagnostic Provider registered attributes and registered tests” on page 189. The next step is for the MBean to actually implement these methods. The methods to implement include:

- “getRegisteredDiagnostics”
- “getDiagnosticProviderName”
- “getDiagnosticProviderID”
- “configDump”
- “stateDump” on page 196
- “selfDiagnostic” on page 196
- “localize” on page 197

getRegisteredDiagnostics

This method exposes the registration information for this Diagnostic Provider. It is commonly used by the DP Utility in the administration console to gather information about Diagnostic Providers that are to be displayed in the console. This method returns a **DiagnosticProviderInfo** object that is usually attained by passing the appropriate XML to a **DiagnosticProviderHelper** helper class. Here is an example:

```
public DiagnosticProviderInfo getRegisteredDiagnostics() {
    InputStream regIS= Thread.currentThread().getContextClassLoader().getResourceAsStream(
        "com/ibm/ws/xxx/SampleDP2DiagnosticProvider.xml");
    dpInfo = DiagnosticProviderHelper.loadRegistry(regIS, sDPName) ;

    if (dpInfo == null) {
        sSampleDP2MBeanLogger.logp(Level.WARNING, sThisClass, "getRegisteredDiagnostics",
            "RasDiag.DPInfo.NoGotz") ;
    }
    return dpInfo ;
}
```

Notice that the XML is packaged and available in the *classpath* of the current *classloader*. The “Registration XML” on page 197 contains crucial information that the Diagnostic Provider uses to “Populate the payload” on page 197 and “localize” on page 197 results.

getDiagnosticProviderName

This is usually a pretty simple return of a constant as the following example shows

```
public String getDiagnosticProviderName() {
    return sDPName;
}
```

getDiagnosticProviderID

This is usually a pretty simple return of a Java Management Extensions (JMX) object ID that MBeans can pull out of the base class method. For example:

```
public String getDiagnosticProviderId() {
    return getObjectname().toString() ;
}
```

configDump

The *configDump* method enables the Diagnostic Provider to expose the configuration data that was in place when this Diagnostic Provider started (or the current values of them). The **DiagnosticEvent** objects that this method returns include a “Payload” on page 197 that contains the core data. The following is an excerpt from a configDump method:

```
public DiagnosticEvent [] configDump(String aAttributeIdSpec, boolean aRegisteredOnly) {
    HashMap cdHash = new HashMap(64) ;

    // "Populate the payload" on page 197
```

```

DiagnosticEvent [] diagnosticEvent = new DiagnosticEvent[1] ;
diagnosticEvent[0] = DiagnosticEventFactory.createConfigDump(getObjectName().toString(),
    "ThisClassName", "configDump", cdHash) ;

return diagnosticEvent ;
}

```

This returns an array of **DiagnosticEvent** objects. Normally, *configDump* and *stateDump* return only one object. However, the method accepts an array because on z/OS systems a server can have multiple servants, and aggregation of the output from the servants is stored in the array.

stateDump

The *stateDump* method enables the Diagnostic Provider to expose the current state data, or data about the current operating conditions of the Diagnostic Provider. The data made available can be anything likely to assist a customer, an IBM support person, or automated tooling in analyzing the health of the component and problem determination if there is an issue. The amount of data available is impacted by the State Collection Specification in effect at the time. If the current State Collection Specification involves the collection of additional data by the Diagnostic Provider, then this additional data can be exposed in the *stateDump*. The **DiagnosticEvent** objects that this method returns include a “Payload” on page 197 that contains the core data. The following is an excerpt from a *stateDump* method:

```

public DiagnosticEvent [] stateDump(String aAttributeIdSpec, boolean aRegisteredOnly) {
    HashMap sdHash = new HashMap(64) ;

    // "Populate the payload" on page 197

    DiagnosticEvent [] diagnosticEvent = new DiagnosticEvent[1] ;
    diagnosticEvent[0] = DiagnosticEventFactory.createStateDump(getObjectName().toString(),
        "ThisClassName", "stateDump", sdHash) ;

    return diagnosticEvent ;
}

```

This returns an array of **DiagnosticEvent** objects. Normally, *configDump* and *stateDump* return only one object.

selfDiagnostic

The *selfDiagnostic* method enables the Diagnostic Provider to perform certain predefined activities to test key functionalities of your system. These tests should not have a lasting effect on the system. For example, if the test is to create a TCP/IP connection to a remote host, the test should also break that connection before returning its results so that the state of the component is unchanged by the test. The information returned by the test is determined by the attributes included in the test section of the XML file. The following is an excerpt from a *selfDiagnostic* method:

```

public DiagnosticEvent [] selfDiagnostic(String aAttributeIdSpec, boolean aRegisteredOnly) {
    TestInfo [] testInfo = dpInfo.selfDiagnosticInfo.testInfo ; // Retrieve the test registry information
    Pattern testChecker = Pattern.compile(aAttributeIdSpec) ; // Compile test regexp parm for faster checking
    ArrayList deList = new ArrayList(8) ; // Allocate expandable list of DiagnosticEvents
    for (int i = 0; i < testInfo.length; i++) {
        if (testChecker.matcher(testInfo[i].id).matches()) {
            HashMap deHash = new HashMap(32) ;

            // "Populate the payload" on page 197

            deList.add(DiagnosticEventFactory.createDiagnosticEvent(getObjectName().toString(),
                DiagnosticEvent.EVENT_TYPE_SELF_DIAGNOSTIC, DiagnosticEvent.LEVEL_INFO,
                "ThisClassName", "selfDiagnostic", dpInfo.resourceBundleName,
                "RasDiag.SDP2.createDE3", // MsgKey for localization
                // Parns to incorporate in msg
                new Object [] { "OneParm", "TwoParm", "RedParm", "BlueParm"}, deHash)) ;

```

```

    }
}

DiagnosticEvent [] diagnosticEvent = new DiagnosticEvent[deList.size()] ;
diagnosticEvent = (DiagnosticEvent [])deList.toArray(diagnosticEvent) ;

return diagnosticEvent ;
}

```

This returns an array of **DiagnosticEvent** objects. In this example, one **DiagnosticEvent** was created from each test that matched the parameter regular expression. The Diagnostic Provider is not required to produce only one per test. The generation of “Payload” is similar to that of *configDump* and *stateDump*.

localize

The **DiagnosticEvents** that methods return contain payload **HashMaps** that contain **MessageKeys** and **ResourceBundles**. The final consumer of these events is often not on the server, and thus may not have the appropriate *classpath* to resolve this. For this purpose, a callback to the Diagnostic Provider to localize the variables is done. A helper method, however, makes it a simple method to write, as this example demonstrates:

```

public String [] localize(String [] aKeys, Locale aLocale) {
    return DiagnosticProviderHelper.localize(dpInfo.resourceBundleName, aKeys, aLocale) ;
}

```

Note that the **dpInfo** (*DiagnosticProviderInfo*) object is needed as this object includes a reference to the **ResourceBundle**.

Payload

A recurring theme in these methods is the ability to include a payload in return objects. This is a set of *name=value* pairs that include the information being exposed by the method. Diagnostic Events returned from a *configDump*, *stateDump*, or *selfDiagnostic* test are relatively complex Java objects. The majority of the information that is returned is contained in the **DiagnosticData** portion of the **DiagnosticEvent** object. Each attribute returned by the Diagnostic Provider is stored in an entry in a **HashMap**. There can be cascading **HashMaps** within a single **DiagnosticEvent** object (if breaking the data down into subGroups makes sense). Each **HashMap** entry contains either a reference to a child **HashMap**, or a **DiagnosticTypedValue** (which contains the value, the type of data, and a **MsgKey** for localization of the label or /name). The values to be returned should be filtered with:

- The type of method (that is, *configDump*, *stateDump*, or *selfDiagnostic*)
- The **AttributeIdSpec** sent in to filter the values
- The current State Collection Specification (which can impact the amount of data available).

Populate the payload

The API documentation for *DiagnosticProviderHelper.queryMatchingDPInfoAttributes* explains how to do the filtering before retrieving the data. In some cases, it is easier and helps performance for a Diagnostic Provider to retrieve all data into the Payload and then filter the **HashMap** after the fact. The post-population filtering can be done with the method *DiagnosticProviderHelper.filterEventPayload*. For information on use of the JavaBeans type approach, see the API documentation for the *AttributeBeanInfo.populateMap* method.

Registration XML

Registration XML enables much of the information needed by the Diagnostic Provider to be externalized. It also provides a means of commonizing localization and consumption of the tests (thus aiding automation). An excerpt of this XML from a sample Diagnostic Provider follows:

```

<!DOCTYPE diagnosticProvider PUBLIC "RasDiag" "/DiagnosticProvider.dtd">

<diagnosticProvider>
  <resourceBundleName> com.ibm.ws.rasdiag.resources.RasDiagSample</resourceBundleName>
  <state>
    <attribute>
      <id>Leg-Foot</id>
      <descriptionKey>SampleDiagnostic.LegFoot.descriptionKey</descriptionKey>
      <registered>true</registered>
    </attribute>
    <attribute>
      <id>Leg-Ankle</id>
      <descriptionKey>SampleDiagnostic.LegAnkle.descriptionKey</descriptionKey>
      <registered>true</registered>
    </attribute>
  </state>
  <config>
    <attribute>
      <id>Arm-Hand-Size</id>
      <descriptionKey>SampleDiagnostic.HandSize.descriptionKey</descriptionKey>
      <registered>true</registered>
    </attribute>
    <attribute>
      <id>Leg-Foot-Size</id>
      <descriptionKey>SampleDiagnostic.FootSize.descriptionKey</descriptionKey>
      <registered>true</registered>
    </attribute>
  </config>
  <selfDiagnostic>
    <test>
      <id>Kick</id>
      <descriptionKey>SampleDiagnostic.Kick.descriptionKey</descriptionKey>
      <attribute>
        <id>Kick-Pain</id>
        <descriptionKey>SampleDiagnostic.KickPain.descriptionKey</descriptionKey>
      </attribute>
      <attribute>
        <id>Kick-Length</id>
        <descriptionKey>SampleDiagnostic.KickLength.descriptionKey</descriptionKey>
      </attribute>
    </test>
    <test>
      <id>Throw</id>
      <descriptionKey>SampleDiagnostic.Throw.descriptionKey</descriptionKey>
      <attribute>
        <id>Throw-Pain</id>
        <descriptionKey>SampleDiagnostic.ThrowPain.descriptionKey</descriptionKey>
        <registered>true</registered>
      </attribute>
      <attribute>
        <id>Throw-Length</id>
        <descriptionKey>SampleDiagnostic.ThrowLength.descriptionKey</descriptionKey>
        <registered>true</registered>
      </attribute>
    </test>
  </selfDiagnostic>
</diagnosticProvider>

```

For understanding the storage of this information into a **DiagnosticProviderInfo** object, see the API documentation for *DiagnosticProviderInfo*. For conceptual information about the purpose of the registration XML, see “Diagnostic Provider registered attributes and registered tests” on page 189.

Diagnostic Provider XML example:

Here is an example of the Diagnostic Provider Extensible Markup Language (XML).


```

version="6.0"
platform="common"
aggregationHandlerClass="com.ibm.ws.management.component.DiagnosticProviderAggregator"
description="DiagnosticProvider portion of MBean for inclusion into MBeans implementing this interface">
  <attribute
    description="DiagnosticProviderName (not dependent on runtime, but subset of ObjectName"
    getMethod="getDiagnosticProviderName" name="diagnosticProviderName"
    type="java.lang.String" proxyInvokeType="unicall" proxySetterInvokeType="multicall"/>
  <operation
    description="Get the DiagnosticProvider ID"
    impact="INFO" name="getDiagnosticProviderId" role="operation"
    targetObjectType="objectReference" type="java.lang.String" proxyInvokeType="unicall">
    <signature/>
  </operation>
  <operation
    description="Return the registry information based on type (config/state/selfDiag)."
    impact="INFO" name="getRegisteredDiagnostics" role="operation"
    targetObjectType="objectReference"
    type="com.ibm.wsspi.rasdiag.diagnosticProviderRegistration.DiagnosticProviderInfo"
    proxyInvokeType="unicall">
    <signature/>
  </operation>
  <operation
    description="Dump the configuration information associated with managed resource."
    impact="INFO" name="configDump" role="operation"
    targetObjectType="objectReference" type="[Lcom.ibm.wsspi.rasdiag.DiagnosticEvent;"
    proxyInvokeType="multicall">
    <signature>
      <parameter description="Attribute ID to use"
        name="attributeId" type="java.lang.String"/>
      <parameter description="Report on just registered info, or all info"
        name="registeredOnly" type="boolean"/>
    </signature>
  </operation>
  <operation
    description="Dump state information for the managed resource."
    impact="INFO" name="stateDump" role="operation"
    targetObjectType="objectReference" type="[Lcom.ibm.wsspi.rasdiag.DiagnosticEvent;"
    proxyInvokeType="multicall">
    <signature>
      <parameter description="Attribute ID to use"
        name="attributeId" type="java.lang.String"/>
      <parameter description="Report on just registered info, or all info"
        name="registeredOnly" type="boolean"/>
    </signature>
  </operation>
  <operation
    description="Perform diagnostics on the managed resource driven by current diagnostic mode setting."
    impact="ACTION" name="selfDiagnostic" role="operation"
    targetObjectType="objectReference" type="[Lcom.ibm.wsspi.rasdiag.DiagnosticEvent;"
    proxyInvokeType="multicall">
    <signature>
      <parameter description="Test ID to use"
        name="testId" type="java.lang.String"/>
      <parameter description="Report on just registered info, or all info"
        name="registeredOnly" type="boolean"/>
    </signature>
  </operation>
  <operation
    description="localize messages for console display"
    impact="INFO" name="localize" role="operation"
    targetObjectType="objectReference" type="[Ljava.lang.String;"
    proxyInvokeType="unicall">
    <signature>

```

```
<parameter description="Message Keys" name="msgKeys" type="[Ljava.lang.String;"/>
<parameter description="Locale to use for output" name="locale" type="java.util.Locale"/>
</signature>
</operation>
```

Creating a Diagnostic Provider registration XML file

The Diagnostic Provider registration XML is used to provide information about the exposed configuration, state, and self diagnostic attributes and tests to the Diagnostic Provider utility. It is also used to populate objects needed later in the process, to assist in filtering, and to assist in localization.

Before you begin

Programming knowledge of your system and the proper authorities to perform the following steps.

About this task

The steps that follow outline a general process for creating a Diagnostic Provider (DP) registration Extensible Markup Language (XML) file.

Procedure

1. Start with the DP document type definition (DTD). If you are using the helper methods (see the step called *Create your DP implementation* in “Creating a Diagnostic Provider” on page 191), you can use this DOCTYPE line to pick up the common DTD:

```
<!DOCTYPE diagnosticProvider PUBLIC "RasDiag" "/DiagnosticProvider.dtd">
```

If you are extending an existing MBean with an existing XML configuration, you might need either to add the DP XML to an existing DTD, or omit the DP XML entirely. If you omit the DP XML, you will not be able to validate that your XML file is well formed.

2. Follow the conventions described in “Diagnostic Provider Extensible Markup Language” on page 192 to help keep your XML consistent with other components. You can find an example of a small DP registration XML file in “Diagnostic Provider method implementation” on page 194.

Associating a Diagnostic Provider ID with a logger

If you are using a Diagnostic Provider to manage alerts and messages, you need to associate the Diagnostic Provider ID with a logger. This can be done dynamically or through a static assignment.

About this task

Components whose diagnostics are managed through a Diagnostic Provider MBean should include the Diagnostic Provider ID (DPID) in all logged messages. In some cases a single logger always logs with the same DPID. In those cases, it is appropriate to statically associate the DPID with the logger. In other cases, a logger might log on behalf of various diagnostic domains. For example, although every data source has a separate Diagnostic Provider MBean, they all share the same logger. In those cases, the DPID can be dynamically supplied on each logging call.

Static Assignment

About this task

The method below statically assigns a DPID to a logger.

Procedure

Associate a DPID with a logger:

```
Logger logger = Logger.getLogger("com.ibm.ws.MyClass");
DiagnosticProviderHelper.addDiagnosticProviderIDtoLogger(logger, dpid);
```

Dynamic Assignment

About this task

DPIDs can be associated with a single log request by including them as the first message parameter, prefixed with **DPID:.** To associate a DPID with a single log request using a logger:

```
Object[] parms = new Object[] { "DPID:" + dpid };
logger.logp(classname, methodname, "MSG0001", parms);
```

Note that in the dynamic case, the DPID does not need to actually show up in the formatted message. The two examples below illustrate:

```
(in resource bundle)
// by not including {0} first parm is not printed in the message.
MSG0001=This message does not include the DPID.
```

```
// note - it is not recommended to print the DPID in your message.
MSG0002=This message includes the DPID...it's value is {0}.
```

It is recommended that messages not include the DPID in the formatted message. As shown above, this is done by not including {0} in the message value in the resource bundle.

Using Diagnostic Providers from wsadmin scripts

In addition to enabling Diagnostic Providers (DP) from the administration console, you can also use them through scripts from the Wsadmin tool.

About this task

You might want to enable, disable, or configure Diagnostic Providers from the administrative console, but in some cases it might be more efficient or useful to do so with scripts using the wsadmin tool.

Read the wsadmin tool information about using the tool with scripts.

Procedure

1. List the MBeans that implement the Diagnostic Provider (DP) interface. Enter

```
$AdminControl queryNames diagnosticProvider=true,*
```

And you will see an output that displays all of the Diagnostic Providers in a format like this:

```
"WebSphere:name=Default DataSource,process=server1,platform=dynamicproxy,node=
camelhair,JDBCProvider=Derby JDBC Provider,
diagnosticProvider=true,j2eeType=JDBCDataSource,J2EEServer=server1,Server=server1,
version=6.1.0.0,type=DataSource,
mbeanIdentifier=cells/camelhairCell/nodes/camelhair/servers/server1/resources.xml#
DataSource_1131113688564,
JDBCResource=Derby JDBC Provider,cell=camelhairCell"
"WebSphere:name=DefaultEJBTimerDataSource,process=server1,platform=dynamicproxy,
node=camelhair,
JDBCProvider=Derby JDBC Provider (XA),diagnosticProvider=true,j2eeType=
JDBCDataSource,J2EEServer=server1,Server=server1,version=6.1.0.0,type=DataSource,
mbeanIdentifier=cells/camelhairCell/nodes/camelhair/servers/server1/
resources.xml#DataSource_1000001,
JDBCResource=Derby JDBC Provider (XA),cell=camelhairCell"
WebSphere:name=WebcontainerDiagnosticProvider,process=server1,platform=
dynamicproxy,node=camelhair,diagnosticProvider=true,
version=6.1.0.0,type=WebcontainerEventProvider,mbeanIdentifier=null,
cell=camelhairCell
```

2. Capture the ObjectName of your Diagnostic Provider in a variable. This enables you to reference your Diagnostic Provider more easily, especially in a script. For example, instead of typing all of those lines, if you want to work with the WebContainer Diagnostic Provider, for example, you can do the following:

- set DP [lindex [\$AdminControl queryNames name=WebcontainerDiagnosticProvider,diagnosticProvider=true,*] 0]

This ObjectName stored in the DP variable can be used on the methods, or you can use the Diagnostic Provider name as text or a variable.

- Now that you have the ObjectName in a variable, you can get the Diagnostic Provider name in a variable with the command:

```
set DPNm [$AdminControl invoke $DS getDiagnosticProviderNameById $DP]
```

This provides the result:

```
WebContainerDP
```

Now the DiagnosticProvider (WebContainer) is addressable by its objectname in variable DP, or by its DiagnosticProvider name in variable DPNm. If you would prefer, you can hard-code the DPName *WebContainerDP* as it is short enough.

3. Save the ObjectName of the DiagnosticService MBean to a variable. For wsadmin, WebSphere Application Server provides this MBean so that the output of the Diagnostic Provider is more easily consumable. Enter

```
set DS [lindex [$AdminControl queryNames name=DiagnosticService,*] 0]
```

4. Run a configDump. You can run a configDump and capture all attributes with the command:

```
$AdminControl invoke $DS configDumpFormattedById [list $DP .* true null]
```

This lists the values that the Diagnostic Provider used at start up (and possible current values). .

Table 38. An excerpt of the configDump output. The following table lists the values that the Diagnostic Provider used at start up and possible current values.

Item Concatenated Name	Value
customProperties =	Null
defaultVirtualHostName =	default_host
jvmProps =	Null
localeProps =	Null
servletCachingEnabled =	false
aliases =	*:9080;*:80;*:9443;

5. Filter the output of your configDump. You can use configDumpFormatted (leaving off the ById) and switch **\$DP** for **\$DPNm** or the string **WebContainerDP**. This example uses *\$DPNm* on this slightly modified version whereby it only picks up attributes dealing with automation:

```
$AdminControl invoke $DS configDumpFormatted [list $DPNm .*auto.* true null]
```

This results in just those attributes that contain **auto** in them. Full (but strict) regular expression syntax is allowed.

Table 39. Results. The following table lists the concatenated names and values.

Item Concatenated Name	Value
autoLoadFiltersEnabled =	false
autoRequestEncoding =	false
autoResponseEncoding =	false
autoLoadFiltersEnabled =	false
autoRequestEncoding =	false
autoResponseEncoding =	false

The syntax is the same for stateDumps and selfDiagnostics

Viewing the run time configuration of a component using Diagnostic Providers

You can use the administrative console to navigate to configuration data that can be used to check the health of a server runtime component.

Before you begin

You must have sufficient authority to run the action.

About this task

Runtime components that have associated diagnostic providers can include their Diagnostic Provider ID (DPID) in their log entries. If you know the DPID, you can enter it directly in the quick link text box. Otherwise, navigate to the desired process by using the tree view displayed at the bottom of the panel, as shown in the steps below.

Procedure

1. Start the administration console.
2. From the task bar on the left side of the console, select **Troubleshooting**.
3. From the task bar on the left side of the console, select **Diagnostic Provider**.
4. From the task bar on the left side of the console, select **Configuration Data**.
5. Either directly enter a Diagnostic Provider ID in the **Quick link using diagnostic provider ID** text box, or select a process (cluster / node / server) from the available processes displayed at the bottom of the panel under the section title **Server selection topology**.
6. From the list of available diagnostic providers for the selected process, choose the desired diagnostic provider name. The configuration data for that diagnostic provider appears.

Configuration data quick link or server selection

Use this panel to select a Diagnostic Provider server for viewing run time configuration data.

To view this administrative console page, click **Troubleshooting > Diagnostic Provider > Configuration Data**

Quick link using Diagnostic Provider ID

From the Configuration data panel, enter a Diagnostic Provider ID to go directly to the page for the configuration data for the Diagnostic Provider for the specific server.

Server selection topology

Use these folders to select server or cluster for viewing the configuration data for a Diagnostic Provider.

If you choose a cluster, whatever action you choose is performed on *each* server in the cluster.

The enterprise applications section shows you the servers that a particular application is running on. If you select a server from this list, the action is performed on that server, not specifically that application.

Diagnostic Providers (selection)

Use this panel to select a Diagnostic Provider from the selected server or cluster.

The list will contain only Diagnostic Providers registered on the selected server or cluster. Not all Diagnostic Providers register with every server in the cell.

You can follow several navigation paths to view this administrative console page. For example, click **Troubleshooting** then expand **Diagnostic Provider** and click on **Tests**. Under **Server selection topology**, click on a server or cluster name then click on a Diagnostic Provider from the list.

Diagnostic Providers

Choose a diagnostic provider from this list.

The path you chose to get to this panel determines which panel displays next.

- If you chose Troubleshooting > Diagnostic Provider > Tests, you see a panel that lists all of the available tests to run on the Diagnostic Provider.
- If you chose Troubleshooting > Diagnostic Provider > State Data, you see a panel that shows the collected state data for the Diagnostic Provider.
- If you chose Troubleshooting > Diagnostic Provider > Configuration Data, you see a panel that shows the configuration data for the Diagnostic Provider.

Configuration data

Use this panel to view the current configuration data for a Diagnostic Provider on a selected server or cluster. Not necessarily every piece of configuration data appears, but data that can be helpful in problem determination is shown.

You can follow several navigation paths to view this administrative console page. For example, click **Troubleshooting** then expand **Diagnostic Provider** and click on **Configuration data** . Under **Server selection topology**, click on a server or cluster name then click on a Diagnostic Provider from the list.

The attributes show information that has been configured for the Diagnostic Provider. You can use the **Save** button to save the information to a file.

Note: Results from a configuration dump contain names that start with either *startup* or *current*. The *startup* entries represent data that was read in by the component at server startup time. The *current* entries contain data that is current – meaning the value of the attributes in use by the runtime at the time the configuration dump was requested.

Node

This is the node name from where the configuration data was collected.

Server

This is the server name from where the configuration data was collected.

Name

This is the name of the attribute for the configuration data.

Value

This is the value of the configuration data.

Description

This is a description of the configuration data.

Viewing the run time state data or configuring the state data collection specifications for a Diagnostic Provider

Use the administrative console to navigate to the state data that can be used to check the health of a server runtime component, or you can configure the state data to be collected for a server.

Before you begin

You must have sufficient authority to execute the action.

About this task

In the server selection topology section, use the view state data radio button to go to the list of registered diagnostic providers. Use the change state data collection specification radio button to modify the state collection specification for the runtime components for a server. Runtime components that have associated diagnostic providers can include their Diagnostic Provider ID (DPID) in their log entries. If you know the DPID, you can enter it directly in the quick link text box.

Procedure

1. Start the administration console.
2. Select **Troubleshooting**.
3. Select **Diagnostic Provider**.
4. Select **State Data**.
5. Select the **View State Data** radio button to simply look at the state data, or select the **Change state data collection specification** radio button to change the configuration.
6. Either directly enter a Diagnostic Provider ID in the **Quick link using diagnostic provider ID** text box, or select a process (cluster / node / server) from the available processes displayed at the bottom of the panel.
 - If you chose the **View State Data** radio button, a panel listing the available Diagnostic Providers appears. Choose one of the providers by clicking on it. A panel displaying the state data appears.
 - If you chose the **Change state data collection specification** radio button, a panel appears that contains a list of the available Diagnostic Providers and a text entry block. The state collection specification for the selected process is managed from this panel. Select one of the available providers by using the checkbox next to it.

Diagnostic Provider State Collection Specification

The State Collection specification provides a mechanism for indicating what additional data diagnostic providers in the system should retain in cases where this additional data could be useful for problem determination or application tuning.

In normal operation, most components should work optimally and not store any operational data that is not needed. There are times, however, when an administrator or automated tool may want a component to collect more information than normal to help in problem determination. This data could then be exposed through a State dump. The State Collection specification was created as a syntax for indicating what additional data the diagnostic providers in the system should retain.

For the syntax of the `aCollectionSpec` string, refer to the `DiagnosticConfigHome` API documentation. It is basically a semicolon (;) separated list of collection specification clauses which are of the form:

```
<DiagnosticProviderName regexp>:<AttributeId regexp>=[0|1]
```

Where the `DiagnosticProviderName` regular expression will make this clause apply to any Diagnostic Provider Name that matches that regular expression. The `AttributeId` regexp and the boolean value (0 for off, and 1 for on) are stored in the `DiagnosticConfig` object that each Diagnostic Provider uses. Turning on or off, and processing the clauses left to right allows relatively complex specification. Any specification that is not explicitly turned on is considered to be off. This format is explained further in the following examples.

To turn on tracing for all attributes in the MyDP Diagnostic Provider:

```
MyDP:.*=1
```

To turn on tracing for all attributes of all Diagnostic Providers (this will probably impact system performance):

```
.*:.*=1
```

To turn on all tracing for all attributes of all Diagnostic Providers beginning with ConnMgr (for example, Data Sources):

```
ConnMgr.*:.*=1
```

This specification turns on special collection attributes in the MyDP Diagnostic Provider that begin with the string PoolInfo. If, however, the attribute begins with PoolInfo.Db2Pool, then the collection is off (because it is read left to right).

```
MyDP:PoolInfo.*=1;MyDP:PoolInfo.Db2Pool.*=0
```

It should be noted that State dumps can return important information even in the case where there is no State Collection Specification turned on for a Diagnostic Provider. Diagnostic Providers frequently have to keep some state information in order to operate. Anything in this category is available in a State dump even if there is no special data collection going on. Using the State Collection Specification may increase the amount of data available.

State Data Quick Link or Server Selection

Use this panel to select a server or cluster to either view collected state data, or to configure state data to collect for a Diagnostic Provider.

To view this administrative console page, click **Troubleshooting > Diagnostic Provider > State Data**.

Quick link using Diagnostic Provider ID

Enter a Diagnostic Provider ID to go directly to the view page for the collected state data for the Diagnostic Provider.

Server selection topology

Use these radio buttons and folders to select a specific server or cluster for viewing of state data or configuring the specification of state data.

If you choose a cluster, whatever action you choose is performed on *each* server in the cluster.

The enterprise applications section shows you the servers that a particular application is running on. If you select a server from this list, the action is performed on that server, not specifically that application.

View state data

Select this radio button to view the state data for a Diagnostic Provider. Then select the cell or cluster you want to work with.

Change state data collection settings

Select this radio button to configure the state collection specification for a Diagnostic Provider. Then select the cluster or managed server you want to work with.

State data

Use this panel to view the current state data for a Diagnostic Provider on a selected server or cluster.

To view this administrative console page, click **Troubleshooting > Diagnostic Provider > State data >** select the View state data radio button and then select a server or cluster name > select a Diagnostic Provider from the list.

The attributes show information that has been collected as part of the enabled state collection specification for the Diagnostic Provider. You can use the **save...** button to save the information to a file

Node

This is the node name from where the state data was collected.

Server

This is the server name from where the state data was collected.

Name

This is the name of the state collection specification used to collect the state data.

Value

This is the value of the state collection specification used to collect the state data.

Description

This is a description of the state collection specification used to collect the state data.

Detailed state specification

Use this panel to view the attributes and descriptions of the Diagnostic Provider that you have selected.

To add attributes, select the checkbox next to your chosen diagnostic provider, then select the **Add to specification** button.

To remove a diagnostic provider's sub-component attribute from the state specification, select the sub-component attribute in the displayed list and then select the **Remove from specification** button.

When you are done adding or removing a diagnostic provider's sub-component attributes, select the **Done** button.

To view this administrative console page, click **Troubleshooting > Diagnostic Provider > State data** > select the View state data radio button and then select a server or cluster name > select a Diagnostic Provider from the list.

Attribute

This is the individual state collection specification available for the Diagnostic Provider.

Description

This is the description of the individual state collection specification item.

Change state specification

Use this panel to add a Diagnostic Provider and its attributes to the specification for collecting state data.

To add a diagnostic provider (DP) and *all* of its attributes, select the checkbox next to your chosen DP, then click on the **Add to specification** button. To add only *some* of the DP's attributes, click on the DP name itself in the list, and a new panel where you can perform this task appears.

To put the state specification into affect, select the **Apply** or **OK** button.

To reset the specification to its original state, use the **Reset** button.

To manually enter a state specification, update the text area with the state specification and use the **Update** button.

To view this administrative console page, click **Troubleshooting > Diagnostic Provider > State data** > select the Change state data collection specification radio button and then select a server or cluster name > select a Diagnostic Provider from the list.

Name

This is a list of available Diagnostic Providers for the server selected.

Modifying the State Collection Specification from wsadmin scripts

In addition to modifying the State Collection Specification from the administrative console, you can also modify these settings using scripts and the wsadmin tool.

About this task

In doing problem determination, you might want to begin collecting additional data during normal processing. This can be accomplished by modifying the State Collection Specification dynamically. This section illustrates how to do that through the wsadmin tool . This technique can be used to turn on traces, as well as to turn off traces. Depending on the usage pattern of the component, the impact should take affect shortly after it is set.

Procedure

1. Capture the DiagnosticService ObjectName into a variable. Enter

```
set DS [lindex [$AdminControl queryNames name=DiagnosticService,*] 0]
```
2. Use this variable to drive the method to set the specification. Enter

```
$AdminControl invoke $DS setStateCollectionSpec "SampleDiagnosticProvider:player.*=1;  
SampleDiagnosticProvider:defense.*=1"
```

The specification is of the form **DiagnosticProviderName:Attributeld=0|1...** (with a semicolon at the end, multiple sub-specifications can be entered similar to the TraceSpec). The DiagnosticProviderName and Attributeld can be proper regular expressions.

Running a self diagnostic on a Diagnostic Provider

You can check the status of server runtime components with predefined tests that can be associated with a Diagnostic Provider. Use the administrative console to access these functions.

Before you begin

You must have sufficient authority to execute the action.

About this task

You can access a list of predefined diagnostic tests that you can use to check the status of a server runtime component. Runtime components that have associated diagnostic providers can include their Diagnostic Provide ID (DPID) in their log entries. If you know the DPID, you can enter it directly in the quick link text box. Otherwise, navigate to the desired process by using the tree view displayed at the bottom of the panel.

Procedure

1. Start the administration console.
2. Select **Troubleshooting**.
3. Select **Diagnostic Provider**.
4. Select **Tests** .
5. Either directly enter a Diagnostic Provider ID in the **Quick link using diagnostic provider ID** text box, or select a process (cluster / node / server) from the available processes displayed in the **Server selection topology** section.
6. Select the desired self diagnostic test.
7. Read the output messages from the self diagnostic test.

8. Select a self diagnostic test message by clicking on it. The console displays a panel with the attributes related to the message you chose.

Tests Quick Link or Server Selection

Use this panel to select a Diagnostic Provider server for diagnostic tests.

To view this administrative console page, click **Troubleshooting > Diagnostic Provider > Tests**.

Quick link using Diagnostic Provider ID

Enter a Diagnostic Provider ID to go directly to the list of tests for the Diagnostic Provider.

Server selection topology

Use these folders to select server or cluster for viewing the available tests for a Diagnostic Provider.

If you choose a cluster, whatever action you choose is performed on *each* server in the cluster.

The enterprise applications section shows you the servers that a particular application is running on. If you select a server from this list, the action is performed on that server, not specifically that application.

Test selection

Use this panel to select one of the tests that are available for the chosen Diagnostic Provider on the chosen server or cluster.

You can follow several navigation paths to view this administrative console page. For example, click **Troubleshooting > Diagnostic Provider > Tests > select a server or cluster name > select a Diagnostic Provider** from the list.

Test identification

Choosing a test ID causes the test to run. Results of the test are shown on the Test Results panel.

Test description

A description of the test available to run on the Diagnostic Provider.

Test Results

Use this panel to see the results from the server or cluster members for the selected test.

You can follow several navigation paths to view this administrative console page. For example, click **Troubleshooting > Diagnostic Provider > Tests > select a cluster name > select a Diagnostic Provider** from the list > select a Test identification from the list.

Multiple results can be returned from a test from each server. The results are sorted by Node, then by Server, then by Severity. You can page through the messages that are returned.

Server

The name of the server where the test result came back from.

Node

The name of the node where the test result came back from.

Severity

The severity of the result from the test run.

Message

A description of the test result.

The entries in this column are linked to another panel. If you click on a message, you can see additional attributes associated with the message.

Test result details

Use this panel to see additional attributes for the selected test result.

To view this administrative console page, click **Troubleshooting** > **Diagnostic Provider** > **Tests** > select a cluster name > select a Diagnostic Provider from the list > select a Test identification from the list > select a message.

The attributes show information that helped to diagnose the condition described in the message. You can use the **Save** button to save to a file the attributes and the messages to which they correspond.

Name

The name of the test.

Value

This is the value of the test result.

Description

This is a description of the test.

Chapter 16. Troubleshooting help from IBM

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

Purchase of WebSphere Application Server entitles you to one year of telephone support under the Passport Advantage® program. For details on the Passport Advantage program, visit http://www.lotus.com/services/passport.nsf/WebDocs/Passport_Advantage_Home.

If you cannot resolve your problem by any of the preceding methods:

1. Go to the WebSphere Application Server support page.
2. Expand the Contact Support section.
 - Click "Exchanging information with IBM Tech Support", and follow the information collection instructions provided.
 - Click "Directory of worldwide contacts" to find the appropriate contact information for your geography.

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.

IBM Support Assistant

Use the IBM Support Assistant to organize your problem solving investigations and get the latest troubleshooting tools to analyze everything from logs to memory dumps. The IBM Support Assistant provides access to many of the same troubleshooting tools that our IBM Technical Support teams use when solving problems.

IBM Support Assistant Data Collector

The IBM Support Assistant Data Collector is a utility for running automated, symptom-specific data collection scripts. These scripts gather logs, trace, configuration, system-specific information, and other symptom-specific data that is useful to diagnose various potential problems. You can attach the resultant collections to an IBM Service Request so that you can get help from IBM Support. The IBM Support Assistant Data Collector is included with WebSphere Application Server, and can be run from your `$app_server_root/bin` directory. IBM Support might ask you to run this tool and submit the output.

Consulting

For complex issues such as integration with legacy systems, education, and help in getting started quickly with the WebSphere product family, consider using IBM consulting services. To learn about these services, browse the website <http://www.ibm.com/services/fullservice.html>.


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.
- *Solve a problem* links take you to specific problems and resolutions documented by WebSphere Application Server technical support personnel.
- The *Download* links provide free WebSphere Application Server maintenance upgrades and problem determination tools.

 The fix packs and refresh packs for the WebSphere Application Server for IBM i products are also provided in the group PTF for the WebSphere Application Server product. The WebSphere Application Server group PTF also includes other IBM i group PTFs which are necessary for running your application servers. For more information see the PTFs page.

- *fixes* are software patches which address specific WebSphere Application Server defects. Selecting a specific defect from the list in the *Fixes by version* page takes you to a description of what problem the fix addresses.
- Fix packs are bundles of multiple fixes, tested together and released as a maintenance upgrade to WebSphere Application Server. Refresh packs are fix packs that also contain new function. 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 *fix list* link on that page displays a list of the fixes which the fix pack includes. If you intend to install a 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.

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 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 Must gather documents for information to gather to send to IBM Support.

Using IBM Support Assistant

IBM Support Assistant is a free troubleshooting application that helps you research, analyze, and resolve problems using various support features and tools. IBM Support Assistant enables you to find solutions yourself using the same troubleshooting techniques used by the IBM Support team, and it allows you to organize and transfer your troubleshooting efforts between members of your team or to IBM for further support.

About this task

IBM Support Assistant helps you diagnose problems with applications that are deployed to WebSphere Application Server. You can track and organize troubleshooting activity, automate the collection of MustGather files, and use tools that analyze artifacts to understand the root cause of errors.

You can customize IBM Support Assistant by installing the problem determination tools that you need to troubleshoot various problems. For example, when the JVM runs out of memory, application threads hang, the JVM crashes, or if performance declines, you can use analysis tools to diagnose these problems.

Procedure

- Follow the installation instructions on IBM Support Assistant (ISA) website at: IBM Support Assistant.
- Learn about suggested add-on tools for IBM Support Assistant that are recommended for use with WebSphere Application Server at Primary Troubleshooting Tools for WebSphere Application Server

Results

Using the IBM Support Assistant Data Collector

The IBM Support Assistant Data Collector for WebSphere Application Server is a tool you can run to gather data from your application server system for problem determination purposes. It replaces the collector tool, which is deprecated.

About this task

Note: The IBM Support Assistant Data Collector for WebSphere Application Server tool focuses on automatic collection of problem data. It also provides symptom analysis support for the various categories of problems encountered by IBM software products. Information pertinent to a type of problem is collected to help identify the origin of the problem under investigation. The tool assists customers by reducing the amount of time it takes to reproduce a problem with the proper RAS tracing levels set, as well as by reducing the effort required to send the appropriate log information to IBM Support.

Procedure

1. Start the tool. The tool runs in console mode by starting the launch script from the command line.
 - In a Windows environment, run the `app_server_root/bin/isadc.bat` command.
 - In a Linux, AIX, HP-UX, Solaris, IBM i, or zOS environment, run the `app_server_root/bin/isadc.sh` command.


Note: You can optionally run the tool from a profile bin directory, instead of the `app_server_root\bin\` directory: `profile_root/bin/isadc.bat` or `profile_root/bin/isadc.sh`.

2. Run the IBM Support Assistant Data Collector tool with the user ID for which you configured your WebSphere Server instance. Depending on what collection you are running, you are asked for additional information to complete the data collection activities. A script might require additional configuration information, information about the sequence of events leading up to the problem you are dealing with, or for your preferences regarding how it completes the collection.


At each step, the choices are presented as numbered lists and you input the number of your selection and press the enter key. When input is required, prompts are displayed at which you enter your response and press the enter key. You can find collection details for each WebSphere Application Server problem type in their corresponding MustGather documents.

The tool has a silent collection capability for recording your responses from a console mode session in a file and uses the file to drive subsequent instances of the same collection script. When running in this mode, you are taken to an ordinary interactive session, where you supply the responses to the script prompts. In addition to influencing the current collection, however, your responses are also saved in the file that you named. When the interactive session completes, you can use this response file to start the same script in the future without the need for explicit user input.

Use the following command to create a response file that contains the answers to all the questions for a certain run through the data collector:

-  `app_server_root/bin/isadc.sh -record response_filename`

Use the following command to provide the response file when starting the tool:

-  `app_server_root/bin/isadc.sh -silent response_filename`

The response file is a plain text file, so you can edit it to change the responses as needed. The file looks very much like a Java properties file, with comments that start with #, and a series of key-value pairs. You can add pauses to a response file using one of the following two keys:

- `PauseScriptTime=X`, where *X* is a positive integer representing the number of seconds a script pauses. If anything other than a positive integer is found, an error message is written to both the console and the log, followed by a message telling the user to hit the enter key when they are ready for the script to proceed.
- `PauseScript=user-defined message` - this string is printed to the console, along with a message to hit the enter key you are ready for the script to proceed.

When using response files, remember that sensitive information, such as user names and passwords, might be stored in these files. It is important that you manage these files in a manner that prevents unauthorized access to sensitive information. Note that passwords are not encrypted.

Tip: If you are not able to provide root or administrator access to the user to run the collection scripts, make sure that the user ID has administrator privileges for your WebSphere Server, for example, `startServer`, `stopServer`, and `wsadmin` commands.

By default, the version of the tool (and the various subcomponents) is printed to the console from which it was launched.

3. Stop the collector tool by typing the quit option in console mode.

Chapter 17. Default behavior for OutOfMemory exceptions

WebSphere Application Server Version 8 Fix Pack 2 and later ships with IBM Java 6 R26 on supported operating systems. Beginning with this version, the default behavior for OutOfMemory (OOM) exceptions has changed.

By default in IBM Java 5 and later, the first four OOM exceptions for the lifetime of a Java process produce a PHD-formatted heap dump, a Java dump file (javacore), and a snap dump file. By default in IBM Java 6 R26 and later, the first OOM for the lifetime of a Java process produces a PHD-formatted heap dump, a Java dump file, a snap dump file, and an operating system dump—core file on Linux, AIX, and IBM i (not to be confused with a javacore), user-mode minidump with full memory on Windows operating systems, and SYSTDUMP on z/OS. The second, third, and fourth OOM exceptions produce only a PHD-formatted heap dump and a Java dump file. Therefore, the change in default behavior is an additional system dump on the first OOM exception.

A system dump is a superset of a PHD heap dump. A system dump also includes memory contents (strings, primitives, variable names, and so on), thread and frame local information, some native memory information, and more. For more information, go to <http://www.ibm.com/developerworks/opensource/library/j-memoryanalyzer/index.html>. This added information can solve a larger class of problems, provide more general insight into a running JVM, and ultimately reduce the time that it takes to solve a problem. In IBM Java 5 R12 or earlier and IBM Java 6 R9 or earlier, a system dump had to be post-processed using the jextract tool. With recent versions of IBM Java, however, including the one that ships with this new OOM behavior, a system dump can be directly loaded by a DTFJ-capable tool such as the Memory Analyzer Tool without any post-processing for OOMs (just like a PHD heap dump). For crashes, however, jextract should still be used.

The default configuration can be observed with the `-Xdump:what` generic JVM argument. For example, running `$WAS/java/bin/java -version -Xdump:what` produces the following output (some output removed). Note the range option in particular.

```
# java -version -Xdump:what
Registered dump agents
...
-Xdump:system:
  events=systhrow,
  filter=java/lang/OutOfMemoryError,
  label=... core.&Y&m&d.&H% M% S.% pid.% seq.dmp,
  range=1..1,
  priority=999,
  request=exclusive+compact+prewalk
...
-Xdump:heap:
  events=systhrow,
  filter=java/lang/OutOfMemoryError,
  label=... heapdump.% Y% m% d.% H% M% S.% pid.% seq.phd,
  range=1..4,
  priority=500,
  request=exclusive+compact+prewalk,
  opts=PHD
...
-Xdump:java:
  events=systhrow,
  filter=java/lang/OutOfMemoryError,
  label=... javacore.% Y% m% d.% H% M% S.% pid.% seq.txt,
  range=1..4,
  priority=400,
  request=exclusive+preempt
...
-Xdump:snap:
  events=systhrow,
```

```
filter=java/lang/OutOfMemoryError,  
label=... Snap.% Y% m% d.% H% M% S.% pid.% seq.trc,  
range=1..4,  
priority=300,  
request=serial
```

Note: A blank space was added to this example after each occurrence of the % character to avoid inappropriate conversions of the text in this information center.

This configuration can be changed using the `-Xdump` (http://publib.boulder.ibm.com/infocenter/javasdk/v6r0/topic/com.ibm.java.doc.diagnostics.60/diag/tools/dumpagents_syntax.html) generic JVM arguments (<http://www-01.ibm.com/support/docview.wss?uid=swg21417365>):

- To revert to the old behavior on OOM, use:
`-Xdump:system:none -Xdump:system:events=gpf+abort+traceassert+corruptcache`
- To remove PHD heap dumps on OOM, use:
`-Xdump:heap:none`
- To remove PHD heap dumps on OOM and take system dumps on not just the first OOM but on the first four:
`-Xdump:heap:none
-Xdump:system:none
-Xdump:system:events=gpf+abort+traceassert+corruptcache
-Xdump:system:events=systhrow,filter=java/lang/OutOfMemoryError,range=1..4,request=exclusive+compact+prewalk`

It is critical that users ensure that operating systems are correctly configured to produce untruncated system dumps:

- **AIX:**
http://publib.boulder.ibm.com/infocenter/javasdk/v6r0/topic/com.ibm.java.doc.diagnostics.60/diag/problem_determination/aix_setup_full_core.html
- **IBM i:**
http://publib.boulder.ibm.com/infocenter/javasdk/v6r0/topic/com.ibm.java.doc.diagnostics.60/diag/problem_determination/i5os_setup_os.html
- **Linux:**
http://publib.boulder.ibm.com/infocenter/javasdk/v6r0/topic/com.ibm.java.doc.diagnostics.60/diag/problem_determination/linux_setup.html
- **z/OS:**
http://publib.boulder.ibm.com/infocenter/javasdk/v6r0/topic/com.ibm.java.doc.diagnostics.60/diag/problem_determination/zos_setup_dumps.htm

System dumps are written to the location specified by the file parameter in `-Xdump` (see the previous label attribute in `-Xdump:what`). You should change this to a dedicated partition with sufficient space. For example:

```
-Xdump:system:file=/var/dumps/core.% Y% m% d.% H% M% S.% pid.% seq.dmp
```

Note: A blank space was added to this example after each occurrence of the % character to avoid inappropriate conversions of the text in this information center.

You might also consider using the `-Xdiagnosticscollector` generic JVM argument, which will print a warning to `native_stderr.log` if the JVM detects that ulimits are set incorrectly on startup. For more information, go to <http://publib.boulder.ibm.com/infocenter/javasdk/v6r0/topic/com.ibm.java.doc.diagnostics.60/diag/tools/verifyingyourjavadiagnosticsconfiguration.html>

Chapter 18. Configuring the memory leak policy

The leak detection policy for the WebSphere® Application Server is turned off by default. You can configure a leak detection, prevention, and action policy to accommodate your applications and environment so that potential memory leaks are reported and acted upon. Leak detection, prevention, and proactive fixing provides for protection and resiliency for servers that face persistent out of memory errors. When a classloader memory leak is detected, WebSphere Application Server notifies you with informational messages in the log and by taking JVM heapdumps so that you can troubleshoot the problem. Optionally, you might also choose to have WebSphere Application Server mitigate, and if possible, fix the memory leak using reflection and other techniques.

Before you begin

A common error in Java Platform, Enterprise Edition (Java EE) applications is a classloader memory leak. A memory leak can result from subtle application bugs or a more complex use, such as third party libraries. System resources, such as CPU time due to garbage collection and the Java heap are consumed when a leak is present. A system can become unresponsive even though all other resources are available. Unless a protection and early warning system is built in, the system might remain in this degraded state and ultimately die due to an out of memory error.

Note: You can configure WebSphere Application Server to detect, prevent, and take action, if possible, on classloader memory leaks using the memory leak detection policy. For more information about memory leaks, read the Memory Leaks in Java Platform, Enterprise Edition applications topic.

About this task

The leak detection option is disabled by default. You can use Java virtual machine (JVM) custom properties to adjust the leak policy values, such as, enable and disable leak detection, action, and prevention. These custom properties are only applicable to a stand-alone server or managed application server and not to a node agent, admin agent, job manager, or deployment manager. When the application or the server is shutting down, WebSphere Application Server determines the classloaders that have leaked and are held references to all of associated loaded classes and objects. If a classloader leak is detected, a heapdump or systemdump is taken.

All persistent configurations of this service are completed using JVM custom properties. There are no administrative console panels. At runtime, use the `MemoryLeakConfig` and `MemoryLeakAdmin` mbeans for configuration and administration respectively; however, the configuration changes are not persisted until JVM custom properties are configured.

The `MemoryLeak` service and its mbean are active only in an application server that hosts applications and services requests. This service is not active on a Deployment Manager, node agent, administrative agent, or other servers types like WebSphere proxy server and so on.

Tip: Configure the JVM `com.ibm.ws.runtime.component.disableMemoryLeakService` custom property to permanently disable the service.

Procedure

1. Create or modify JVM custom properties to enable various aspects of the memory leak service, based on the table that follows. Read Java virtual machine custom properties to change the values of the JVM custom properties.

Table 40. `com.ibm.ws.runtime.component.MemoryLeakConfig.detectAppCLLeaks`

Information	Value
Name	<code>com.ibm.ws.runtime.component.MemoryLeakConfig.detectAppCLLeaks</code>

Table 40. *com.ibm.ws.runtime.component.MemoryLeakConfig.detectAppCLLeaks* (continued)

Information	Value
Description	When the server is shutting down or an application stops, WebSphere Application Server determines the classloaders that have leaked and issues warnings and other additional information that aids in debugging the memory leak. See also the Improved classloader leak detection PMR.
Default	false

Table 41. *com.ibm.ws.runtime.component.MemoryLeakConfig.clearAppCLLeaks*

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.clearAppCLLeaks
Description	Enable proactive classloader leak mediation and fixing. When this property is set to true, WebSphere Application Server mediates on behalf of the application to remedy any classloader leaks that are detected.
Default	false

Table 42. *com.ibm.ws.runtime.component.MemoryLeakConfig.preventJreMemoryLeaks*

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.preventJreMemoryLeaks
Description	Enable WebSphere Application Server to eliminate certain classes of memory leaks that are caused by the JRE loading singletons on the thread context classloader.
Default	true

Table 43. *com.ibm.ws.runtime.component.MemoryLeakConfig.generateHeapDumps*

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.generateHeapDumps
Description	Set to true to cause a heapdump to be created when a memory leak is detected .
Default	true

Table 44. *com.ibm.ws.runtime.component.MemoryLeakConfig.leakSweeperDelay*

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.leakSweeperDelay
Description	Delay after an application or module stops to check for classloader leaks.
Default	10000 (milliseconds)

Table 45. *com.ibm.ws.runtime.component.MemoryLeakConfig.monitorSystemApps*

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.monitorSystemApps
Description	Set to true to watch for memory leaks on system applications that are shipped by IBM, such as the administrative console.
Default	false

Attention: You can fine tune the leak action policy behavior using the following thread and timer leaks JVM custom properties. The ThreadLocal, thread, timer and static leak custom properties are applicable only if `com.ibm.ws.runtime.component.MemoryLeakConfig.clearAppCLLeaks` is set to true.

Table 46. *com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesInterruptThreads*. Thread and timer leaks

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesInterruptThreads
Description	Set to true for WebSphere Application Server to attempt to interrupt threads that are started by the web application. Interrupting threads is performed using the <code>Thread.interrupt()</code> method. There is a possibility that the target thread might not respond to the interrupt.
Default	true

Table 47. *com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesHttpClientKeepAliveThread*. Thread and timer leaks

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesHttpClientKeepAliveThread
Description	If an HttpClient keep-alive timer thread is not started by this web application and is still running, WebSphere Application Server changes the context class loader from the current classloader to its parent to prevent a memory leak. The keep-alive timer thread stops on its own when the keep-alive threads all die. However, on a busy system that might not happen for some time.
Default	true

Table 48. *com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesStopTimerThreads*. Thread and timer leaks

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesStopTimerThreads
Description	Set to true for WebSphere Application Server to stop any java.util.TimerThreads that are started by the web application.
Default	false

Table 49. *com.ibm.ws.runtime.component.MemoryLeakConfig.jvmThreadGroupNames*. Thread and timer leaks

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.jvmThreadGroupNames
Description	List of ThreadGroup names to ignore when scanning for threads that are started by the web application that must be shut down. This list is delineated by underscores.
Default	system_RMI Runtime

Attention: The following static class variable leaks properties apply only if `com.ibm.ws.runtime.component.MemoryLeakConfig.clearAppCLLeaks` is set to true.

Table 50. *com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesStatic*. Static class variable leaks

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesStatic
Description	Set to true for WebSphere Application Server to attempt to make final fields null from loaded classes when a web application stops, as a work around for garbage collection bugs and application coding errors. Applications without memory leaks using recent JVMs should operate correctly when this option is set to false.
Default	false

Table 51. *com.ibm.ws.runtime.component.MemoryLeakConfig.filterPrefixes*. Static class variable leaks

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.filterPrefixes
Description	Member attributes starting with these filters are not set to null when <code>clearReferencesStatic</code> is true.
Default	java javax com.ibm org sun com.sun. The list is delineated by spaces.

Attention: The following Threadlocal leaks properties apply only if `com.ibm.ws.runtime.component.MemoryLeakConfig.clearAppCLLeaks` is set to true.



Table 52. *com.ibm.ws.runtime.component.MemoryLeakConfig.checkThreadLocalLeaks*. Threadlocal leaks

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.checkThreadLocalLeaks
Description	Determines if ThreadLocal leaks should be checked when an application stops.
Default	false



Table 53. *com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesThreadLocal*. Threadlocal leaks

Information	Value
Name	com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesThreadLocal

Table 53. *com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesThreadLocal* (continued). Threadlocal leaks

Information	Value
Description	Eliminate threadlocal leaks by renewing threads in the threadpool when an application stops. This property applies only to distributed platforms.
Default	true

Attention: The following properties apply only if `clearReferencesThreadLocal` is set to true.

IBM i

Table 54. *com.ibm.ws.runtime.component.MemoryLeakConfig.renewThreadPoolNames*. Threadlocal leaks

Information	Value
Name	<code>com.ibm.ws.runtime.component.MemoryLeakConfig.renewThreadPoolNames</code>
Description	Threadpool names that WebSphere Application Server must renew when a threadlocal leak is detected.
Default	"WebContainer" list is delineated by spaces.

IBM i

Table 55. *com.ibm.ws.runtime.component.MemoryLeakConfig.threadPoolRenewalDelayFactor*. Threadlocal leaks

Information	Value
Name	<code>com.ibm.ws.runtime.component.MemoryLeakConfig.threadPoolRenewalDelayFactor</code>
Description	Control the total amount of time to wait for thread pool to be renewed. The amount of time to wait for threadpool renewal = <code>threadPoolRenewalDelayFactor * threadPool.getKeepAliveTime()</code>
Default	10. For the web container threadpool, the renewal delay is (10 * 60000 ms) 6 minutes.

IBM i

Table 56. *com.ibm.ws.util.ThreadPool.DEFAULT_THREAD_RENEWAL_DELAY*. Threadlocal leaks

Information	Value
Name	<code>com.ibm.ws.util.ThreadPool.DEFAULT_THREAD_RENEWAL_DELAY</code>
Description	Threads in the pool renew after an application stops. Use this custom property to avoid renewing all threads at the same time. This delay is observed between two threads being renewed.
Default	1000 (value is in ms). If the value is negative, threads are never renewed.

These JVM custom properties are persisted in the WebSphere Application Server configuration model in the `server.xml` file. The following code snippet displays the persisted leak policy configuration from the `server_home/config/cells/nodes/servers/server.xml` file of an unmanaged server:

```
<jvmEntries xmi:id="JavaVirtualMachine_1183122130078" verboseModeClass="true" verboseModeGarbageCollection="true" verboseModeJNI="false"
runHProf="false" hprofArguments="" debugMode="false" debugArgs="-agentlib:jdwp=transport=dt_socket,server=y,suspend=n,address=7777"
genericJvmArguments="-agentlib:getClasses -Xquickstart -Xalwaysclassgc" executableJarFileName="" disableJIT="false">
<systemProperties xmi:id="Property_1317048628648" name="com.ibm.ws.runtime.component.MemoryLeakConfig.detectAppCLLeaks" value="true" />
<systemProperties xmi:id="Property_1318975518491" name="com.ibm.ws.runtime.component.MemoryLeakConfig.clearAppCLLeaks" value="true" />
<systemProperties xmi:id="Property_1318955284241" name="com.ibm.ws.runtime.component.MemoryLeakConfig.generateSystemDumps" value="false" />
<systemProperties xmi:id="Property_1319119976147" name="com.ibm.ws.runtime.component.MemoryLeakConfig.generateHeapDumps" value="true" />
<systemProperties xmi:id="Property_1317048628649" name="com.ibm.ws.runtime.component.MemoryLeakConfig.monitorSystemApps" value="false" />
</jvmEntries>
```

2. Click **Apply**.
3. Click **OK**.
4. Save the changes.
5. Restart the Application Server for the changes to take effect.

Example

The Memory Leak policy for WebSphere Application Server can be configured and persisted using JVM custom properties as described in this example. At runtime the memory leak detection, prevention and policy configuration can be changed using the `MemoryLeakConfig` mbean.

Administration of the Memory leak policy can be carried out using the `MemoryLeakAdmin` mbean. The leak policy affects how the application server responds to a classloader memory leak when an application or when the server is stopped.

You can adjust the Memory leak policy settings by using the wsadmin scripting interface. These changes take effect immediately, but do not persist to the server configuration, and are lost when the server is restarted. The following script provides an example of how to configure and administer the memory leak policy using wsadmin jacl scripting :

```
# Scripting in JACL

# Get the object name of the MemoryLeak Configuration object you want to change the values on
set leakConfig [$AdminControl completeObjectName "type=MemoryLeakConfig,*"]
WebSphere:cell=smitaNode03Cell,name=LeakConfig,type=MemoryLeakConfig,node=smitaNode03,process=server1

# Get the object name of the MemoryLeak Administration object you want to issue operations
set leakAdmin [$AdminControl completeObjectName "type=MemoryLeakAdmin,*"]
WebSphere:cell=smitaNode03Cell,name=LeakAdmin,type=MemoryLeakAdmin,node=smitaNode03,process=server1

# Look at all the attributes of the MemoryLeakConfig mbean
wsadmin>$Help all $leakConfig
Name: WebSphere:cell=smitaNode03Cell,name=LeakConfig,type=MemoryLeakConfig,node=smitaNode03,process=server1
Description: Information on the management interface of the MBean
Class name: com.ibm.ws.runtime.component.MemoryLeakConfig

Attribute                                     Type                                           Access
JvmThreadGroupNames                         java.lang.String                             RW
FilterPrefixes                              java.lang.String                             RW
RenewThreadPoolNames                        java.lang.String                             RW
DetectAppCLLeaks                            boolean                                       RW
ClearAppCLLeaks                             boolean                                       RW
MonitorSystemApps                           boolean                                       RW
NoDumps                                     boolean                                       RW
GenerateHeapDumps                           boolean                                       RW
GenerateSystemDumps                         boolean                                       RW
ClearReferencesStatic                       boolean                                       RW
ClearReferencesInterruptThreads             boolean                                       RW
ClearReferencesStopTimerThreads             boolean                                       RW
ClearReferencesHttpClientKeepAliveThread    boolean                                       RW
ClearReferencesThreadLocal                  boolean                                       RW
LeakSweeperDelay                            int                                           RW
ThreadPoolRenewalDelayFactor                 int                                           RW
PreventJreMemoryLeaks                       boolean                                       RW
LeakConfiguration                           java.lang.String                             RO

Operation
Notifications
Constructors

# Print the current Memory Leak Policy configuration on the console
wsadmin>$AdminControl getAttribute $leakConfig LeakConfiguration
MemoryLeakConfig [getClass()=class com.ibm.ws.runtime.component.MemoryLeakConfig, hashCode()=37266644
preventJreMemoryLeaks true
detectAppCLLeaks true
    monitorSystemApps false
    leakSweeperDelay 10000
clearAppCLLeaks true
    clearReferencesStopTimerThreads false
    clearReferencesHttpClientKeepAliveThread true
    clearReferencesInterruptThreads true
        jvmThreadGroupNames [system, RMI Runtime]
    clearReferencesStatic true
        filterPrefixes [java., javax., com.ibm., org., sun., com.sun]
    clearReferencesThreadLocal true
        renewThreadPoolNames [WebContainer]
        threadPoolRenewalDelayFactor 1
noDumps false
    generateHeapDumps true
    generateSystemDumps false

# Change the configuration
wsadmin>$AdminControl setAttribute $leakConfig ThreadPoolRenewalDelayFactor 10
wsadmin>$AdminControl setAttribute $leakConfig ClearReferencesStopTimerThreads true

# See the updated configuration
```

```

wsadmin>$AdminControl getAttribute $leakConfig LeakConfiguration
MemoryLeakConfig [getClass()=class com.ibm.ws.runtime.component.MemoryLeakConfig, hashCode()=37266644
preventJreMemoryLeaks true
detectAppCLLeaks true
    monitorSystemApps false
    leakSweeperDelay 10000
clearAppCLLeaks true
    clearReferencesStopTimerThreads true
    clearReferencesHttpClientKeepAliveThread true
    clearReferencesInterruptThreads true
        jvmThreadGroupNames [system, RMI Runtime]
    clearReferencesStatic true
        filterPrefixes [java., javax., com.ibm., org., sun., com.sun]
    clearReferencesThreadLocal true
        renewThreadPoolNames [WebContainer]
        threadPoolRenewalDelayFactor 10
noDumps false
    generateHeapDumps true
    generateSystemDumps false

# Look at all the operations of the MemoryLeakAdmin mbean
wsadmin>$Help all $leakAdmin
Name: WebSphere:cell=smitaNode03Cell,name=LeakAdmin,type=MemoryLeakAdmin,node=smitaNode03,process=server1
Description: Information on the management interface of the MBean
Class name: com.ibm.ws.runtime.component.MemoryLeakAdmin
Operation
java.lang.String findLeaks()
java.lang.String fixLeaks()
java.lang.String fixLeaks(java.lang.String)

# Find current classloader memory leaks
wsadmin>$AdminControl invoke $leakAdmin findLeaks
CWMLL0028I: The following web applications were stopped (reloaded, undeployed),
but their classes from previous runs are still loaded in memory,
thus causing a memory leak. [[78577.075.724.NWALogging#NWALoggingEJB.jar]].

# Fix ALL current classloader memory leaks
wsadmin>$AdminControl invoke $leakAdmin fixLeaks
CWMLL0036I: Please watch the SystemOut log for results of the fix leak operation.

wsadmin>$AdminControl invoke $leakAdmin fixLeaks {"78577.075.724.NWALogging#NWALoggingEJB.jar"}
CWMLL0036I: Please watch the SystemOut log for results of the fix leak operation.

```

Memory leaks in Java Platform, Enterprise Edition applications

Memory leaks come in various types such as thread and ThreadLocal leaks, ClassLoader leaks, system resource leaks, and connection leaks. Approaches to memory leak detection typically involve examination of Java virtual machine tool Interface (JVMTI) or performance monitoring infrastructure (PMI) counters to watch for slow growth in Java or native heap usage.

Note: WebSphere Application Server Version 8.5 provides a top down pattern-based memory leak detection, prevention, and action by watching for suspect patterns in application code at run time. WebSphere Application Server has some means of protection against memory leaks when stopping or redeploying applications. If leak detection, prevention and action are enabled, WebSphere Application Server monitors application and module activity performs diagnostic actions to detect and fix leaks when an application or an individual module stops. This feature helps in increasing application up time with frequent application redeployments without cycling the server.

Class loader memory leak

Many memory leaks manifest themselves as class loader leaks. A Java class is uniquely identified by its name and the class loader that loaded it. Classes with the same name can be loaded multiple times in a single JVM, each in a different class loader. Each web application gets its own class loader and this is what WebSphere Application Server uses for isolating applications.

An object retains a reference to the class it is an instance of. A class retains a reference to the class loader that loaded it. The class loader retains a reference to every class it loaded. Retaining a reference to a single object from a web application pins every class loaded by the web application. These references often remain after a web application reload. With each reload, more classes are pinned which leads to an out of memory error.

Class loader memory leaks are normally caused by the application code or JRE triggered code.

JRE triggered leaks

Memory leaks occur when Java Runtime Environment (JRE) code uses the context class loader to load an application singleton. These singletons can be threads or other objects that are loaded by the JRE using the context class loader.

If the web application code triggers the initialization of a singleton or a static initializer, the following conditions apply:

- The context class loader becomes the web application class loader.
- A reference is created to the web application class loader. This reference is never garbage collected.
- Pins the class loader, and all the classes loaded by it, in memory.

Application triggered leaks

Categories of application triggered leaks are as follows:

- Custom ThreadLocal class
- Web application class instance as ThreadLocal value
- Web application class instance indirectly held through a ThreadLocal value
- ThreadLocal pseudo-leak
- ContextClassLoader and threads created by web applications
- ContextClassLoader and threads created by classes loaded by the common class loader
- Static class variables
- JDBC driver registration: RMI targets

For more information on application triggered links, see <http://wasdynacache.blogspot.com/2012/01/websphere-classloader-memory-leak.html>, http://www.websphereusergroup.org.uk/wug/files/presentations/31/Ian_Partridge_-_WUG_classloader_leaks.pdf, and <http://www.ibm.com/support/docview.wss?uid=swg1PM39870>.

WebSphere Application Server now has some means of protection against memory leaks when stopping or redeploying applications. WebSphere Application Server monitors application and module activity and performs diagnostic actions when an application or an individual module is stopped.

There are three parts to this memory leak feature in WebSphere Application Server: detection, prevention, and action.

1. **Detection:** Issue warnings when a memory leak is detected. Through a combination of standard API calls and some reflection tricks when a web application is stopped, undeployed or reloaded. WebSphere Application Server checks for known causes of memory leaks and issues warnings when an application leak is detected, as follows:

```
[11/17/11 12:01:05:911 EST] 00000005 LeakDetection E CWMML0015E: The web application [WasSwat#WasSwatWeb.war] created a ThreadLocal with key of type [test.memleak.MyThreadLocal] (value [test.memleak.MyThreadLocal@216c691]) and a value of type [test.memleak.MyCounter] (value [test.memleak.MyCounter@21942ff]) but failed to remove it when the web application was stopped.
```

2. **Prevention** is on by default and applies only to JRE triggered leaks. JRE triggered leaks are prevented by initializing singletons at server startup, when the application Server class loader is the context class loader.

3. Action: Take proactive action to fix memory leaks. These actions have reasonable defaults and are configured on a case-by-case basis.

```
protected void com.ibm.ws.classloader.clearReferences(){
    if(ENABLE_CLEAR_REFERENCES_JDBC)
        clearReferencesJdbc();
    if(ENABLE_CLEAR_REFERENCES_THREADS)
        clearReferencesThreads();
    if(ENABLE_CLEAR_REFERENCES_THREADLOCALS)
        clearReferencesThreadLocals();
    if(ENABLE_CLEAR_REFERENCES_RMI_TARGETS)
        clearReferencesRmiTargets();
    if(ENABLE_CLEAR_REFERENCES_STATICS)
        clearReferencesStaticFinal();
}
```

Table 57. Leak Fix Summary Matrix

Leak cause	How to fix	WebSphere Application Server Java Virtual Machine properties for enabling and controlling
Threadlocal	Renew threads in the threadpool for a configurable period. Getting threads out of the pool enables the threads and threadlocals to be garbage collected.	<ul style="list-style-type: none"> com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesThreadLocal com.ibm.ws.runtime.component.MemoryLeakConfig.renewThreadPoolNames com.ibm.ws.runtime.component.MemoryLeakConfig.threadPoolRenewalDelayFactor com.ibm.ws.util.ThreadPool.DEFAULT_THREAD_RENEWAL_DELAY
HttpClient keep-alive threads	Switch thread to the parent class loader.	<ul style="list-style-type: none"> com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesHttpClientKeepAliveThread
Timer threads	Use reflection to stop any new tasks that might be scheduled.	<ul style="list-style-type: none"> com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesStopTimerThreads
Non JVM controlled threads	If the thread is started using an executor, shut down the executor or interrupt the thread.	<ul style="list-style-type: none"> com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesInterruptThreads
JDBC drivers	Unregister any JDBC drivers that are registered by the web application that the web application forgot.	<ul style="list-style-type: none"> com.ibm.ws.runtime.component.MemoryLeakConfig.preventJreMemoryLeak
ResourceBundle	Clear the ResourceBundle cache of any bundles that were loaded by this class loader or any class loader where this loader is a parent class loader.	<ul style="list-style-type: none"> com.ibm.ws.runtime.component.MemoryLeakConfig.preventJreMemoryLeak
RMI targets	Use reflection to clear the values in sun.rmi.transport.ObjectTable.implTable and sun.rmi.transport.ObjectTable.objTable.	<ul style="list-style-type: none"> com.ibm.ws.runtime.component.MemoryLeakConfig.preventJreMemoryLeak
Static class variables	WebSphere Application Server nullifies the value of all static class variables of classes that are loaded by the application or module class loader.	<ul style="list-style-type: none"> com.ibm.ws.runtime.component.MemoryLeakConfig.clearReferencesStatic com.ibm.ws.runtime.component.MemoryLeakConfig.filterPrefixes

Chapter 19. Collecting Java dumps and core files using the administrative console

You can use the Java runtime environment to create dump and core files to help with troubleshooting. You can use the administrative console to trigger the creation of these dumps and core files.

About this task

The Java virtual machine (JVM) is capable of producing Java dump and core files to aid in troubleshooting. You can use heap dump and system dump files to help you diagnose memory-related problems, such as memory leaks. You can use Java core files to help you diagnose problems where the CPU is persistently 100% busy, when threads are hanging, or where threads are in a deadlock.

Note: The process of generating dump and core files can have a noticeable performance impact on your system that can last for many seconds or minutes. If necessary, use your test and development environments to better understand the impact of generating dump and core files.

Note: Generating a heap dump or a system dump is not supported for a non-IBM JVM. If you attempt to generate a heap dump or a system dump using the administrative console against a non-IBM JVM, the following message is displayed in the administrative console: oror

Procedure

1. In the navigation pane, click **Troubleshooting > Java dumps and cores**.
2. Select the server or servers you need to collect a dump or core from.
3. Click **System Dump** or **Java Core** or **Heap Dump** depending on your need.

Results

The system dump or heap dump or Java core is created and stored in the profile root directory of the server from which you requested the dump or core.

Java dump and core collection

Use this page to generate various Java dumps and cores from within the administrative console.

To view this administrative console page, click **Troubleshooting > Java dumps and cores**.

Select one or more of the listed servers, and click one of the following buttons to generate dumps. You can analyze dumps and cores using problem determination tools available for the IBM Support Assistant.

Note: Generating a heap dump or a system dump is not supported for a non-IBM Java virtual machine (JVM). If you attempt to generate a heap dump or a system dump using the administrative console against a non-IBM JVM, the following message is displayed in the administrative console: oror

Heap dump

A heap dump is a snapshot of JVM memory. It shows live objects in the memory and references between them.

You can use this option to debug conditions such as memory leaks, heap fragmentation or to investigate what objects are consuming the largest part of the memory.

Java core

Use this button to investigate why a server is hanging or investigate messages in the logs that indicate a thread has not completed its work in the expected amount of time.

System dump

Use this button to generate system native dumps of the server process. These dumps can be quite large.

Chapter 20. Directory conventions

References in product information to *app_server_root*, *profile_root*, and other directories imply specific default directory locations. This article describes the conventions in use for WebSphere Application Server.

Default product locations - IBM i

These file paths are default locations. You can install the product and other components in any directory where you have write access. You can create profiles in any valid directory where you have write access. Multiple installations of WebSphere Application Server products or components require multiple locations.

app_client_root

The default installation root directory for the Application Client for IBM WebSphere Application Server is the `/QIBM/ProdData/WebSphere/AppClient/V85/client` directory.

app_client_user_data_root

The default Application Client for IBM WebSphere Application Server user data root is the `/QIBM/UserData/WebSphere/AppClient/V85/client` directory.

app_client_profile_root

The default Application Client for IBM WebSphere Application Server profile root is the `/QIBM/UserData/WebSphere/AppClient/V85/client/profiles/profile_name` directory.

app_server_root

The default installation root directory for WebSphere Application Server - Express is the `/QIBM/ProdData/WebSphere/AppServer/V85/Express` directory.

java_home

Table 58. Root directories for supported Java Virtual Machines.

This table shows the root directories for all supported Java Virtual Machines (JVMs).

JVM	Directory
32-bit IBM Technology for Java	<code>/QOpenSys/QIBM/ProdData/JavaVM/jdk60/32bit</code>
64-bit IBM Technology for Java	<code>/QOpenSys/QIBM/ProdData/JavaVM/jdk60/64bit</code>

plugins_profile_root

The default Web Server Plug-ins profile root is the `/QIBM/UserData/WebSphere/Plugins/V85/webserver/profiles/profile_name` directory.

plugins_root

The default installation root directory for Web Server Plug-ins is the `/QIBM/ProdData/WebSphere/Plugins/V85/webserver` directory.

plugins_user_data_root

The default Web Server Plug-ins user data root is the `/QIBM/UserData/WebSphere/Plugins/V85/webserver` directory.

product_library

product_lib

This is the product library for the installed product. The product library for each Version 8.5 installation on the system contains the program and service program objects (similar to `.exe`, `.dll`, `.so` objects) for the installed product. The product library name is `QWAS85x` (where `x` is `A`, `B`, `C`, and so on). The product library for the first WebSphere Application Server Version 8.5 product installed on the system is `QWAS85A`. The `app_server_root/properties/product.properties` file contains the value for the product library of the installation, was `.install.library`, and is located under the `app_server_root` directory.

profile_root

The default directory for a profile named *profile_name* for WebSphere Application Server - Express is the `/QIBM/UserData/WebSphere/AppServer/V85/Express/profiles/profile_name` directory.

shared_product_library

The shared product library, which contains all of the objects shared by all installations on the system, is QWAS85. This library contains objects such as the product definition, the subsystem description, the job description, and the job queue.

user_data_root

The default user data directory for WebSphere Application Server - Express is the `/QIBM/UserData/WebSphere/AppServer/V85/Express` directory.

The profiles and profileRegistry subdirectories are created under this directory when you install the product.

The *user_data_root* directory contains the default locations for WLP_USR_DIR and WLP_OUTPUT_DIR when the Liberty profile is installed. These directories are `user_data_root/wlp/usr` and `user_data_root/wlp/output/servers`, respectively.

web_server_root

The default web server path is `/www/web_server_name`.

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.

APACHE INFORMATION. This information may include all or portions of information which IBM obtained under the terms and conditions of the Apache License Version 2.0, January 2004. The information may also consist of voluntary contributions made by many individuals to the Apache Software Foundation. For more information on the Apache Software Foundation, please see <http://www.apache.org>. You may obtain a copy of the Apache License at <http://www.apache.org/licenses/LICENSE-2.0>.

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 Intellectual Property & Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
USA

Trademarks and service marks

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

Microsoft and Windows 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.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

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

Index

A

- application security setting
 - restrictions 7
- Application Server
 - logging 6

D

- data sources
 - restrictions
 - runtime environments 7
- directory
 - installation
 - conventions 227

E

- EAR files
 - restrictions 7
- enterprise bundles
 - restrictions 7

F

- feature restrictions
 - restrictions 7
- FFDC log files
 - logging 6
 - restrictions 7
- FFDC services
 - logging 6

H

- HTTPS port
 - restrictions 7

J

- JDBC data source
 - restrictions
 - runtime environments 7
- JPA annotations
 - restrictions 7

- JSF support
 - restrictions 7

L

- Liberty profile
 - restrictions 7
- log files
 - restrictions 7
- logging configuration 6

O

- OSGi applications
 - runtime environments 7

S

- server configuration files
 - logging 6
- server configurations
 - logging 6
- Server Version
 - logging 6

T

- Troubleshooter reference
 - messages 9

U

- user ID
 - restrictions 7

W

- web applications
 - restrictions 7

X

- XMI files
 - restrictions 7