

IBM SDK for z/OS, Java Technology Edition
Version 6

SDK Guide

IBM

IBM SDK for z/OS, Java Technology Edition
Version 6

SDK Guide

IBM

Note

Before using this information and the product it supports, read the information in "Notices" on page 119.

Copyright information

This edition of the user guide applies to the IBM 31-bit SDK for z/OS, Java Technology Edition, Version 6, product 5655-I98, and to the IBM 64-bit SDK for z/OS, Java Technology Edition, Version 6, product 5655-I99, and to all subsequent releases, modifications, and Service Refreshes, until otherwise indicated in new editions.

Portions © Copyright 1997, 2013, Oracle and/or its affiliates.

© **Copyright IBM Corporation 2003, 2013.**

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

Contents

Preface	v	Securing JAXP processing against malformed input	31
Chapter 1. Overview	1	XML reference information	31
Version compatibility	1	Debugging Java applications.	36
Migrating from earlier IBM SDK or JREs	1	Java Debugger (JDB)	36
Supported environments	2	Selective debugging	37
The z/OS batch toolkit	3	Determining whether your application is running on a 31-bit or 64-bit JVM	38
Chapter 2. Contents of the SDK and Runtime Environment.	5	Determining which JVM version your application is running on	38
Contents of the Runtime Environment	5	How the JVM processes signals.	39
Contents of the SDK.	6	Signals used by the JVM	39
Chapter 3. Installing and configuring the SDK	9	Linking a native code driver to the signal-chaining library.	41
Working with BPXPRM settings	9	Writing JNI applications	42
Setting the region size	10	Supported compilers	43
Setting MEMLIMIT.	10	Native formatting of Java types long, double, float.	43
Setting LE runtime options	10	Support for thread-level recovery of blocked connectors.	44
Setting LE 31-bit runtime options	11	Configuring large page memory allocation	44
Setting LE 64-bit runtime options	12	CORBA support	45
Marking failures.	12	System properties for tracing the ORB	46
Setting the path	12	System properties for tuning the ORB	47
Setting the class path	13	Java security permissions for the ORB	47
Updating your SDK or JRE for daylight saving time changes.	13	ORB implementation classes.	48
Running the JVM under a different code page.	14	RMI over IIOP	48
Using non-default system fonts.	15	Implementing the Connection Handler Pool for RMI	48
Chapter 4. Running Java applications	17	Enhanced BigDecimal	48
The java and javaw commands.	17	Working in a multiple network stack environment	49
Obtaining version information	17	Support for XToolkit	49
Specifying Java options and system properties.	18	Support for the Java Attach API	49
Standard options	19	Chapter 6. Applet Viewer	53
Globalization of the java command	20	Distributing Java applications	53
The Just-In-Time (JIT) compiler.	20	Chapter 7. Class data sharing between JVMs	55
Disabling the JIT	21	Overview of class data sharing	55
Enabling the JIT	21	Class data sharing command-line options	57
Determining whether the JIT is enabled	21	Creating, populating, monitoring, and deleting a cache	61
Specifying a garbage collection policy	22	Performance and memory consumption	62
Garbage collection options	22	Considerations and limitations of using class data sharing.	62
More effective heap usage using compressed references	23	Cache size limits.	62
Pause time	23	JVMTI RetransformClasses() is unsupported	62
Pause time reduction	24	Required APAR for Shared Classes	63
Environments with very full heaps	24	Working with BPXPRMxx settings.	63
Euro symbol support	25	Runtime bytecode modification.	63
Using Indian and Thai input methods	25	Operating system limitations	64
Chapter 5. Developing Java applications	27	Using SharedClassPermission	64
Using XML	27	Adapting custom class loaders to share classes	65
Migrating to the XL-TXE-J	29		

Chapter 8. Service and support for independent software vendors	67
Chapter 9. Accessibility	69
Keyboard traversal of JComboBox components in Swing	69
Chapter 10. Any comments on this user guide?	71
Appendix. Appendixes	73
Command-line options	73
Specifying command-line options	73

General command-line options	74
System property command-line options	75
JVM command-line options	82
JVM -XX command-line options	97
JIT and AOT command-line options	98
Garbage Collector command-line options	101
Default settings for the JVM	111
Known issues and limitations	113
Support for virtualization software	118
Notices	119
Privacy Policy Considerations	121
Trademarks	121

Preface

This user guide provides general information about the IBM® SDK for z/OS®, Java™ Technology Edition, Version 6. The user guide gives specific information about any differences in the IBM implementation compared with the Oracle implementation.

This user guide applies to IBM SDK for z/OS, Java Technology Edition, Version 6.

Read this user guide with the more extensive documentation on the Oracle Web site: <http://www.oracle.com/technetwork/java/index.html>.

Late breaking information about the IBM SDK for z/OS, Java Technology Edition, Version 6 that is not available in the user guide can be found here: <http://www.ibm.com/support/docview.wss?uid=swg21587401>.

The Diagnostics Guide provides more detailed information about the IBM Virtual Machine for Java.



The terms *Runtime Environment* and *Java Virtual Machine* are used interchangeably throughout this user guide.

This user guide is part of a release and is applicable only to that particular release. Make sure that you have the user guide appropriate to the release you are using.

IBM SDK for z/OS, V6 service refresh 9 and earlier

The user guide is available in the code package. Technical changes made for a version of the user guide are indicated by blue chevrons.

IBM SDK for z/OS, V6 service refresh 10 and later

The user guide is available online. Any modifications made to this user guide to support a later service refresh are indicated by graphic images, for example:  In this image, SR11 indicates that changes are made for service refresh 11. End of changes are marked by 

The user guide is also available for downloading as a PDF. Technical changes made for this version are indicated by vertical bars at the beginning of the line.

The Program Code is not designed or intended for use in real-time applications such as (but not limited to) the online control of aircraft, air traffic, aircraft navigation, or aircraft communications; or in the design, construction, operation, or maintenance of any nuclear facility.

Chapter 1. Overview

The IBM SDK is a development environment for writing and running applets and applications that conform to the Java 6 Core Application Program Interface (API).

Version compatibility

In general, any application that ran with a previous version of the SDK should run correctly with the IBM SDK for z/OS, V6. Classes that are compiled with this release are not guaranteed to work on previous releases.

For information about compatibility issues between releases, see the Oracle Web site at:

<http://www.oracle.com/technetwork/java/javase/compatibility-137541.html>

<http://www.oracle.com/technetwork/java/javase/compatibility-137462.html>

<http://www.oracle.com/technetwork/java/javase/compatibility-j2se1-141394.html>

<http://www.oracle.com/technetwork/java/javase/compatibility-135119.html>

If you are using the SDK as part of another product (for example, IBM WebSphere® Application Server), and you upgrade from a previous level of the SDK, perhaps v5.0, serialized classes might not be compatible. However, classes are compatible between service refreshes.

Migrating from earlier IBM SDK or JREs

From Version 5.0, the IBM Runtime Environment for z/OS contains new versions of the IBM Virtual Machine for Java and the Just-In-Time (JIT) compiler.

If you are migrating from an older IBM Runtime Environment, note that:

- The XL TXE-J compiler replaces the XSLT4J interpreter as the default XSLT processor. If you are migrating applications from older versions of Java, see “Migrating to the XL-TXE-J” on page 29.
- The JVM shared library `libjvm.so` is now stored in `jre/lib/<arch>/j9vm` and `jre/lib/<arch>/classic`.
- From Version 5.0 onwards, the JVM Monitoring Interface (JVMMI) is no longer available. You must rewrite JVMMI applications to use the JVM Tool Interface (JVMTI) instead. The JVMTI is not functionally the equivalent of JVMMI. For information about JVMTI, see <http://download.oracle.com/javase/6/docs/technotes/guides/jvmti/> and the Diagnostics Guide.
- From Version 5.0 onwards, the implementation of JNI conforms to the JNI specification, but differs from the Version 1.4.2 implementation. It returns copies of objects rather than pinning the objects. This difference can expose errors in JNI application code. For information about debugging JNI code, see **-Xcheck:jni** in “JVM command-line options” on page 82.
- From Version 5.0 onwards, the format and content of garbage collector verbose logs obtained using **-verbose:gc** have changed. The data is now formatted as XML. The data content reflects the changes to the implementation of garbage

collection in the JVM, and most of the statistics that are output have changed. You must change any programs that process the verbose GC output so that they will work with the new format and data. See the Diagnostics Guide for an example of the new verbose GC data.

- SDK 1.4 versions of the IBM JRE included JVM specific classes in a file called `core.jar`. From Version 5.0 onwards, these are included in a file called `vm.jar`.
- From Version 6, JVM classes are held in multiple JAR files in the `jre/lib` directory. This replaces the single `rt.jar` and `core.jar` from earlier releases.
- For additional industry compatibility information, see Oracle's Java 6 Compatibility Documentation: <http://www.oracle.com/technetwork/java/javase/compatibility-137541.html>
- For additional deprecated API information, see Oracle's Java 6 Deprecated API List: <http://download.oracle.com/javase/6/docs/api/deprecated-list.html>
- All z/OS Java SDK program products can be installed and executed on the same z/OS system. They are independent program products and can coexist in any combination.
- The serial reusability feature of the IBM SDK for z/OS, version 1.4.2 (31-bit) and earlier, started using **-Xresettable**, is not supported. If you specify **-Xresettable** the JVM will issue an error message and will not start. The **-Xinitacsh** and **-Xinitth** options, which allowed heap sizes to be specified for the resettable JVM, are ignored. You can share data between JVMs in an address space (the old **-Xjvmset** and **-Xscmax** options) using Chapter 7, "Class data sharing between JVMs," on page 55, a new facility for Version 5.0. If you specify **-Xjvmset** or **-Xscmax** the JVM will issue an error message and will not start.
- The system property **os.arch** for IBM SDK for z/OS, version 1.4.2 (31-bit) versions and earlier had a value of **390**. From Java 5.0 onwards, the value of **os.arch** is **s390**.
- Tracing class dependencies, started using **-verbose:Xclassdep**, is not supported. If you specify **-verbose:Xclassdep**, the JVM will issue an error message and will not start.
- The JVM detects the operating system locale and sets the language preferences accordingly. For example, if the locale is set to **fr_FR**, JVM messages will be printed in French. To avoid seeing JVM messages in the language of the detected locale, remove the file `$SDK/jre/bin/java_xx.properties` where `xx` is the locale, such as **fr**, and the JVM will print messages in English.

Many new features and capabilities, which might present planning considerations, can be found here: [Summary of changes](#).

Supported environments

The IBM SDK for z/OS, V6 is supported on certain hardware platforms and operating systems, and is tested on specific virtualization environments.

Hardware platform

The z/OS 31-bit and 64-bit SDKs run on System z[®] hardware. The SDKs run on the following servers or equivalents:

- zEnterprise[®] EC12
- z196
- z9-109
- z990

- z900
- z890
- z800

Operating system

The following table shows the latest operating system level tested for each platform architecture. The table indicates whether support for an operating system release was included at the "general availability" (GA) date for the SDK, or at a specific service refresh (SR) level.

Table 1. Supported Operating System levels

Operating system	SDK release supported	Comments
z/OS 1.6	GA	Operating system no longer supported.
z/OS 1.7	GA	Operating system no longer supported.
z/OS 1.8	GA	
z/OS 1.9	GA	
z/OS 1.10	SR3	
z/OS 1.11	SR6	
z/OS 1.12	SR9	
z/OS 1.13	SR10	

Virtualization software

For information about the virtualization software tested for IBM SDK for z/OS, V6, see "Support for virtualization software" on page 118.

The z/OS batch toolkit

From Version 5, Service Refresh 3 onwards, the z/OS products have been enhanced with the JZOS batch toolkit. This toolkit addresses many of the functional and environmental shortcomings in the previous Java batch capabilities on z/OS. The enhancements include a native launcher for running Java applications directly as batch jobs or started tasks and a set of Java methods that make access to traditional z/OS data and key system services directly available from Java applications. Additional system services include console communication, multiline WTO (write to operator), and return code passing capability. For more details, see <http://www.ibm.com/systems/z/os/zos/tools/java/products/jzos/overview.html> and <http://www.ibm.com/systems/z/os/zos/tools/java/>.

Chapter 2. Contents of the SDK and Runtime Environment

The SDK contains several development tools and a Java Runtime Environment (JRE). This section describes the contents of the SDK tools and the Runtime Environment.

Applications written entirely in Java must have **no** dependencies on the IBM SDK's directory structure (or files in those directories). Any dependency on the SDK's directory structure (or the files in those directories) might result in application portability problems.

The , and the accompanying copyright files are the only documentation included in this SDK for z/OS. You can view Oracle's software documentation by visiting the Oracle Web site, or you can download Oracle's software documentation package from the Oracle Web site: <http://www.oracle.com/technetwork/java/index.html> . Additional z/OS related information is available on the z/OS Java Web site at <http://www.ibm.com/systems/z/os/zos/tools/java/>.

Contents of the Runtime Environment

A list of classes, tools, and other files that you can use with the standard Runtime Environment.

- Core Classes - These classes are the compiled class files for the platform and must remain compressed for the compiler and interpreter to access them. Do not modify these classes; instead, create subclasses and override where you need to.
- Trusted root certificates from certificate signing authorities - These certificates are used to validate the identity of signed material. The IBM Runtime Environment for Java contains an expired GTE CyberTrust Certificate for compatibility reasons. This certificate might be removed for later versions of the SDK. See "Expired GTE Cybertrust Certificate" on page 116 for more information.
- JRE tools - The following tools are part of the Runtime Environment and are in the `/usr/lpp/java/J6.0[_64]/jre/bin` directory unless otherwise specified.

ikeycmd (iKeyman command-line utility)

Allows you to manage keys, certificates, and certificate requests from the command line. For more information see the accompanying Security documentation, which includes the *iKeyman User Guide*.

ikeyman (iKeyman GUI utility)

Allows you to manage keys, certificates, and certificate requests. For more information see the accompanying Security documentation, which includes the *iKeyman User Guide*. There is also a command-line version of this utility.

java (Java Interpreter)

Runs Java classes. The Java Interpreter runs programs that are written in the Java programming language.

javaw (Java Interpreter)

Runs Java classes in the same way as the **java** command does, but does not use a console window.

jextract (Dump extractor)

Converts a system-produced dump into a common format that can be used by **jdmpview**. For more information, see **jdmpview**.

keytool (Key and Certificate Management Tool)

Manages a keystore (database) of private keys and their associated X.509 certificate chains that authenticate the corresponding public keys.

kinit

Obtains and caches Kerberos ticket-granting tickets.

klist

Displays entries in the local credentials cache and key table.

ktab

Manages the principal names and service keys stored in a local key table.

pack200

Transforms a JAR file into a compressed pack200 file using the Java gzip compressor.

policytool (Policy File Creation and Management Tool)

Creates and modifies the external policy configuration files that define your installation's Java security policy.

rmid (RMI activation system daemon)

Starts the activation system daemon so that objects can be registered and activated in a Java virtual machine (JVM).

rmiregistry (Java remote object registry)

Creates and starts a remote object registry on the specified port of the current host.

tnameserv (Common Object Request Broker Architecture (CORBA) transient naming service)

Starts the CORBA transient naming service.

unpack200

Transforms a packed file produced by **pack200** into a JAR file.

Contents of the SDK

A list of tools and reference information that is included with the standard SDK.

The following tools are part of the SDK and are located in the `/usr/lpp/java/J6.0[_64]/bin` directory:

appletviewer (Java Applet Viewer)

Tests and runs applets outside a Web browser.

apt (Annotation Processing Tool)

Finds and executes annotation processors based on the annotations present in the set of specified source files being examined.

extcheck (Extcheck utility)

Detects version conflicts between a target jar file and currently-installed extension jar files.

hwkeytool

Manages a keystore of private keys and their associated X.509 certificate chains authenticating the corresponding public keys.

idlj (IDL to Java Compiler)

Generates Java bindings from a given IDL file.

jar (Java Archive Tool)

Combines multiple files into a single Java Archive (JAR) file.

jarsigner (JAR Signing and Verification Tool)

Generates signatures for JAR files and verifies the signatures of signed JAR files.

java (Java Interpreter)

Runs Java classes. The Java Interpreter runs programs that are written in the Java programming language.

java-rmi.cgi (HTTP-to-CGI request forward tool)

Accepts RMI-over-HTTP requests and forwards them to an RMI server listening on any port.

javac (Java Compiler)

Compiles programs that are written in the Java programming language into bytecodes (compiled Java code).

javadoc (Java Documentation Generator)

Generates HTML pages of API documentation from Java source files.

javah (C Header and Stub File Generator)

Enables you to associate native methods with code written in the Java programming language.

javap (Class File Disassembler)

Disassembles compiled files and can print a representation of the bytecodes.

javaw (Java Interpreter)

Runs Java classes in the same way as the `java` command does, but does not use a console window.

jconsole (JConsole Monitoring and Management Tool)

Monitors local and remote JVMs using a GUI. JMX-compliant.

jdumpview (Cross-platform dump formatter)

Analyzes dumps. For more information, see [Using system dumps and the dump viewer](#).

keytool (Key and Certificate Management Tool)

Manages a keystore (database) of private keys and their associated X.509 certificate chains that authenticate the corresponding public keys.

native2ascii (Native-To-ASCII Converter)

Converts a native encoding file to an ASCII file that contains characters encoded in either Latin-1 or Unicode, or both.

policytool (Policy File Creation and Management Tool)

Creates and modifies the external policy configuration files that define your installation's Java security policy.

rmic (Java Remote Method Invocation (RMI) Stub Converter)

Generates stubs, skeletons, and ties for remote objects. Includes RMI over Internet Inter-ORB Protocol (RMI-IIOP) support.

rmid (RMI activation system daemon)

Starts the activation system daemon so that objects can be registered and activated in a Java virtual machine (JVM).

rmiregistry (Java remote object registry)

Creates and starts a remote object registry on the specified port of the current host.

schemagen

Creates a schema file for each namespace referenced in your Java classes.

serialver (Serial Version Command)

Returns the serialVersionUID for one or more classes in a format that is suitable for copying into an evolving class.

tnameserv (Common Object Request Broker Architecture (CORBA) transient naming service)

Starts the CORBA transient naming service.

wsgen

Generates JAX-WS portable artifacts used in JAX-WS Web services.

wsimport

Generates JAX-WS portable artifacts from a Web Services Description Language (WSDL) file.

xjc

Compiles XML Schema files.

Include Files

C headers for JNI programs.

Demos

The demo directory contains a number of subdirectories containing sample source code, demos, applications, and applets that you can use.

copyright

The copyright notice for the SDK for z/OS software.

Chapter 3. Installing and configuring the SDK

See the z/OS Web site for instructions about ordering, downloading, installing, and verifying the SDK.

<http://www.ibm.com/systems/z/os/zos/tools/java/>

Working with BPXPRM settings

Some of the parameters in PARMLIB member **BPXPRMxx** might affect successful Java operation by imposing limits on resources that are required.

The parameters described here do not cover the ones required for Class data sharing. See “Considerations and limitations of using class data sharing” on page 62 for the parameters required for Class data sharing.

To see the current **BPXPRMxx** settings, enter the z/OS operator command `D OMVS,0`. To show the highwater usage for some of the limits, enter the command `D OMVS,L`. If you configure the **BPXPRMxx LIMMSG** parameter to activate the support, **BPXInnnI** messages are reported when the usage approaches and reaches the limits. You can use the `SETOMVS` command to change the settings without requiring an IPL.

Other products might impose their own requirements, but for Java the important parameters and their suggested minimum values are:

Table 2. *BPXPRM settings*

Parameter	Value
MAXPROCSYS	900
MAXPROCUSER	512
MAXUIDS	500
MAXTHREADS	10 000
MAXTHREADTASKS	5 000
MAXASSIZE	2 147 483 647
MAXCPU	2 147 483 647
MAXMMAPAREA	40 960
IPCSEMIDS	500
IPCSEMSEMS	1 000
SHRLIBRGNSIZE	67 108 864
SHRLIBMAXPAGES	4 096

The lower of **MAXTHREADS** and **MAXTHREADTASKS** limits the number of threads that can be created by a Java process.

MAXMMAPAREA limits the number of 4K pages that are available for memory-mapped jar files through the environment variable **JAVA_MMAP_MAXSIZE**.

SHRLIBRGNSIZE controls how much storage is reserved in each address space for mapping shared DLLs that have the +1 extended attribute set. If this storage space

is exceeded, DLLs are loaded into the address space instead of using a single copy of z/OS UNIX System Services storage that is shared between the address spaces. Some of the Java SDK DLLs have the +l extended attribute set. The z/OS command **D OMVS,L** shows the **SHRLIBRGNSIZE** size and peak usage. If this size is set to a much higher value than is needed, Java might have problems acquiring native (z/OS 31-bit) storage. These problems can cause a z/OS abend, such as 878-10, or a Java OutOfMemoryError.

SHRLIBMAXPAGES is only available in z/OS 1.7 and earlier releases. This parameter is like the **SHRLIBRGNSIZE** parameter, except that it is a number of 4K pages. The parameter only applies to DLLs that have the .so suffix, but without the +l extended attribute. This feature requires Extended System Queue Area (ESQA), therefore must be used carefully.

For further information about the use of these parameters, see:

- z/OS MVS™ Initialization and Tuning Reference (SA22-7592)
- z/OS Unix System Services Planning Guide (GA22-7800)

Setting the region size

Java requires a suitable z/OS region size to operate successfully. It is suggested that you do not restrict the region size, but allow Java to use what is necessary. Restricting the region size might cause failures with storage-related error messages or abends such as 878-10.

The region size might be affected by the following factors:

- **JCL REGION** parameter
- **BPXPRMxx MAXASSIZE** parameter
- RACF OMVS segment **ASSIZEMAX** parameter
- IEFUSI

You might want to exclude OMVS from using the IEFUSI exit by setting **SUBSYS(OMVS,NOEXITS)** in PARMLIB member SMFPRMxx.

For further information, see the documentation about the host product under which Java runs.

Setting MEMLIMIT

z/OS uses region sizes to determine the amount of storage available to running programs. For the 64-bit product, set the **MEMLIMIT** parameter to include at least 1024 MB plus the largest expected JVM heap size value **-Xmx**.

See *Limiting Storage use above the bar in z/Architecture* for information about setting the **MEMLIMIT** parameter: <http://www.ibm.com/support/techdocs/atmastr.nsf/WebIndex/FLASH10165>.

Setting LE runtime options

LE runtime options can affect both performance and storage usage. The optimum settings will vary according to the host product and the Java application itself, but it is important to have good general settings.

The LE runtime options are documented in *Language Environment Programming Reference (SA22-7562)* at <http://publibz.boulder.ibm.com/epubs/pdf/ceea3180.pdf>.

Java and other products that are written in C or C++ might have LE runtime options embedded in the main programs by using `#pragma runopts`. These options are chosen to provide suitable default values that assist the performance in a typical environment. Any runtime overrides that you set might alter these values in a way that degrades the performance of Java or the host product. The host product's documentation might provide details of the product's default settings. Changes to the product's `#pragma runopts` might occur as a result of version or release changes. For details of how LE chooses the order of precedence of its runtime options, refer to the *Language Environment Programming Guide (SA22-7561)* at <http://publibz.boulder.ibm.com/epubs/pdf/ceea2180.pdf>.

Use the LE runtime option `RPTOPTS(ON)` as an override to write the options that are in effect, to `stderr` on termination. See the host product documentation and the *Language Environment Programming Guide (SA22-7561)* at <http://publibz.boulder.ibm.com/epubs/pdf/ceea2180.pdf> for details of how to supply LE runtime overrides. Before creating runtime overrides, run the application without overrides, to determine the existing options based on LE defaults and `#pragma` settings.

To tune the options, use the LE runtime option `RPTSTG(ON)` as an override, but be aware that performance could be reduced when you use this option. The output for `RPTSTG(ON)` also goes to `stderr` on termination. The *Language Environment Debugging Guide (GA22-7560)* at <http://publibz.boulder.ibm.com/epubs/pdf/ceea1180.pdf> explains `RPTSTG(ON)` output.

Setting LE 31-bit runtime options

There are a number of LE 31-bit options that are important for successful Java operation.

These options are as follows:

- **ANYHEAP**
- **HEAP**
- **HEAPPOLS**
- **STACK**
- **STORAGE**
- **THREADSTACK**

You can change any, or all, of these options, however if you set the wrong values this might affect performance. The following values are a suggested starting point for these options:

```
ANYHEAP(2M,1M,ANY,FREE)
HEAP(80M,4M,ANY,KEEP)
HEAPPOLS(ON,8,10,32,10,128,10,256,10,1024,10,2048,10,0,10,0,10,0,10,0,10,0,10)
STACK(64K,16K,ANY,KEEP,128K,128K)
STORAGE(NONE,NONE,NONE,0K)
THREADSTACK(OFF,64K,16K,ANY,KEEP,128K,128K)
```

ANYHEAP and **HEAP** initial allocations (parameter 1) might be too large for transaction-based systems such as CICS®. Java applications that use many hundreds of threads might need to adjust the **STACK** initial and increment allocations (parameters 1, 2, 5 and 6) based on the **RPTSTG(ON)** output, which shows the maximum stack sizes that are used by a thread inside the application.

HEAPPOOLS(ON) should improve performance, but the LE-supplied default settings for the cell size and percentage pairs are not optimized for the best performance or storage usage.

For additional information, including how to set the LE runtime options, see:

- the Diagnostics Guide
- the *z/OS Language Environment Programming Reference (SA22-7562)* at <http://publibz.boulder.ibm.com/epubs/pdf/ceea3180.pdf>
- the *z/OS Language Environment Programming Guide (SA22-7561)* at <http://publibz.boulder.ibm.com/epubs/pdf/ceea2180.pdf>
- the host product documentation

Setting LE 64-bit runtime options

There are 64-bit versions of some of the runtime options.

These 64-bit options are as follows:

- **HEAP64**
- **HEAPPOOLS64**
- **STACK64**
- **THREADSTACK64**

A suggested start point for **HEAP64** as an override is **HEAP64(512M,4M,KEEP,16M,4M,KEEP,0K,0K,FREE)**.

The following LE defaults should be appropriate:

```
STACK64(1M,1M,128M) THREADSTACK64(OFF,1M,1M,128M)
HEAPPOOLS64(OFF,8,4000,32,2000,128,700,256,350.1024,100,2048,50,
3072,50,4096,50,8192,25,16384,10,32768,5,65536,5)
```

Before you set an override for **HEAPPOOLS64**, use **RPTOPTS(ON)** or **RPTSTG(ON)** and check the result of `#pragma runopts`. Check this because the host product might have already set cell sizes and numbers that are known to produce good performance.

Also, these settings are dependant on a suitable **MEMLIMIT** setting. Based on these suggested LE 64-bit runtime options, the JVM requirement is a minimum of 512 MB as set for **HEAP64** (which should include **HEAPPOOLS64**), plus an initial value for **STACK64** of 1 MB times the expected maximum number of concurrent threads, plus the largest expected JVM heap **-Xmx** value.

Marking failures

The Java launcher can mark the z/OS Task Control Block (TCB) with an abend code when the launcher fails to load the VM or is terminated by an uncaught exception. To start TCB marking, set the environment variable **IBM_JAVA_ABEND_ON_FAILURE=Y**.

By default, the Java launcher will not mark the TCB.

Setting the path

If you alter the **PATH** environment variable, you will override any existing Java launchers in your path.

About this task

The **PATH** environment variable enables z/OS to find programs and utilities, such as **javac**, **java**, and **javadoc** tool, from any current directory. To display the current value of your **PATH**, type the following command at a command prompt:

```
echo $PATH
```

To add the Java launchers to your path:

1. Edit the shell startup file in your home directory (typically `.bashrc`, depending on your shell) and add the absolute paths to the **PATH** environment variable; for example:

```
export PATH=/usr/lpp/java/J6.0[_64]/bin:/usr/lpp/java/J6.0[_64]/jre/bin:$PATH
```

2. Log on again or run the updated shell script to activate the new **PATH** environment variable.

Results

After setting the path, you can run a tool by typing its name at a command prompt from any directory. For example, to compile the file `Myfile.java`, at a command prompt, type:

```
javac Myfile.java
```

Setting the class path

The class path tells the SDK tools, such as **java**, **javac**, and the **javadoc** tool, where to find the Java class libraries.

About this task

You should set the class path explicitly only if:

- You require a different library or class file, such as one that you develop, and it is not in the current directory.
- You change the location of the `bin` and `lib` directories and they no longer have the same parent directory.
- You plan to develop or run applications using different runtime environments on the same system.

To display the current value of your **CLASSPATH** environment variable, type the following command at a shell prompt:

```
echo $CLASSPATH
```

If you develop and run applications that use different runtime environments, including other versions that you have installed separately, you must set the **CLASSPATH** and **PATH** explicitly for each application. If you run multiple applications simultaneously and use different runtime environments, each application must run in its own shell prompt.

Updating your SDK or JRE for daylight saving time changes

You can apply recent changes to daylight saving time using the IBM Time Zone Update Utility for Java (JTZU).

About this task

Many countries around the world use a daylight saving time (DST) convention. Typically, clocks move forward by one hour during the summer months to create more daylight hours during the afternoon and less during the morning. This practice has many implications, including the need to adjust system clocks in computer systems. Occasionally, countries change their DST start and end dates. These changes can affect the date and time functions in applications, because the original start and end dates are programmed into the operating system and in Java software. To avoid this problem you must update operating systems and Java installations with the new DST information.

The Olson time zone database is an external resource that compiles information about the time zones around the world. This database establishes standard names for time zones, such as "America/New_York", and provides regular updates to time zone information that can be used as reference data. To ensure that Java JREs and SDKs contain up to date DST information, IBM incorporates the latest Olson update into each Java service refresh. To find out which Olson time zone update is included for a particular service refresh, see https://www.ibm.com/developerworks/java/jdk/dst/olson_table.html.

If a DST change has been introduced since the last service refresh, you can use JTZU to directly update your Java installation. You can also use this tool to update your installation if you are unable to move straight to the latest service refresh. JTZU is available from IBM developerWorks® using the following link: <https://www.ibm.com/developerworks/java/jdk/dst/jtzu.html>.

Results

After updating your Java installation with any recent DST changes, your application can handle time and date calculations correctly.

Running the JVM under a different code page

Java for z/OS is shipped in only one version, for EN_US. To run Java under a different locale, convert all the text files in your Java installation from the default (IBM-1047) to your code page.

The **jdkconv** utility converts text files to a local encoding. The utility might be helpful if you get error messages about an invalid format for text files. To check if your code page setting might be the cause of the problem, find which locale you are using by checking the environment variables **LANG** or **LC_ALL**. If the locale value is not C or EN_US then you might see the invalid format message.

To convert your Java installation to a different code page, use the **jdkconv** utility. The utility requires a tool called **cpmod**, which is also provided in your Java installation.

1. Make a copy of your Java installation directory. The reason is that the script overwrites files in the Java installation directory. If you want to undo the changes, you must either reinstall Java, or restore your copy of the directory.
2. The **jdkconv** utility itself is shipped in code page IBM-1047. Before you run the utility, convert it to your code page as follows:

```
cp -p jdkconv jdkconv.backup
iconv -f IBM-1047 -t CODEPAGE jdkconv.backup >jdkconv
chmod 755 jdkconv
```

- In this command sequence, replace **CODEPAGE** with your code page.
3. Ensure that the directory containing both **jdkconv** and **cpmod** is in your **PATH** setting.
 4. Run **jdkconv** as follows:

```
jdkconv CODEPAGE JAVATREE
```

In this command, **CODEPAGE** is your code page, and **JAVATREE** is your Java installation directory.

After running the **jdkconv** utility, test the changes by running a Java application that is sensitive to the system locale.

Using non-default system fonts

If your operating system is z/OS V1.13 or earlier, you must install fonts separately, and optionally edit the Java fonts configuration properties file. For z/OS V2.1 and later, fonts are provided with the operating system by default, and no configuration is required.

About this task

Earlier releases of the z/OS operating system did not include font packages. From V2.1, fonts from the following packages are included by default, in the `/usr/lpp/fonts/worldtype` directory:

- AFP Font Collection for S/390™ (5648-B33)
- IBM Infoprint Fonts for z/OS V1.1 (5648-E76)
- World Type fonts that were not previously available in the z/OS operating system but form part of the InfoPrint Font Collection V3.1

Procedure

Complete these steps if you have z/OS V 1.13 or earlier. Later versions do not require configuration.

1. Purchase the font packages that you require. For example, you could purchase Infoprint WorldType Fonts for AFP Clients as part of program 5648-E77.
2. Install the font packages. Further configuration depends on where you install the packages:
 - If you install into the following directory, you do not have to do any further configuration, but the fonts are removed when you apply a service refresh, and must be replaced afterwards:
 - For IBM SDK Java Technology Edition, V6 service refresh 14 and earlier:
`lib_dir/fonts`
 - For IBM SDK Java Technology Edition, V6 service refresh 15 and later:
`/usr/lpp/fonts/worldtype`
 - If you install to a location outside `install_dir`, the fonts are unaffected when you apply a service refresh. However, you must edit the Java font configuration properties file, which is affected by service refreshes, so changes that you make must be reapplied.

Note: These instructions assume that the font files are called `mtsansdj.ttf` and `tnrwt_j.ttf`. Depending on the font package that you purchased, your font file names might be different, in which case you must rename the files. For an example, see PM05140: Missing fonts on z/OS.

3. If you installed in a directory other than *lib_dir/fonts* or */usr/lpp/fonts/worldtype*, modify the *lib_dir/fontconfig.properties.src* file by changing the paths in the following section:

```
# Font File Names
filename.-Monotype-TimesNewRomanWT-medium-r-normal--*-%d-75-75-*
-*-jisx0208.1983-0=font_path/tnrwt_j.ttf
filename.-Monotype-SansMonoWT-medium-r-normal--*-%d-75-75-*-ji
sx0208.1983-0=font_path/mtsansdj.ttf
```

Where *font_path* depends on your level of IBM SDK Java Technology Edition, V6:

- For IBM SDK Java Technology Edition, V6 service refresh 14 and earlier:
\$JRE_LIB_FONTS
- For IBM SDK Java Technology Edition, V6 service refresh 15 and later:
/usr/lpp/fonts/worldtype

Change the paths to match your font install location, then save a copy of the properties file so you can reapply the changes after you upgrade to a new service refresh.

Related information:

 [Information for Acquiring and Installing Fonts for z/OS Java](#)

Chapter 4. Running Java applications

Java applications can be started using the **java** launcher or through JNI. Settings are passed to a Java application using command-line arguments, environment variables, and properties files.

The **java** and **javaw** commands

An overview of the **java** and **javaw** commands.

Purpose

The **java** and **javaw** tools start a Java application by starting a Java Runtime Environment and loading a specified class.

The **javaw** command is identical to **java**, and is supported on z/OS for compatibility with other platforms.

Usage

The JVM searches for the initial class (and other classes that are used) in three sets of locations: the bootstrap class path, the installed extensions, and the user class path. The arguments that you specify after the class name or `.jar` file name are passed to the main function.

The **java** and **javaw** commands have the following syntax:

```
java [ options ] <class> [ arguments ... ]
java [ options ] -jar <file.jar> [ arguments ... ]
javaw [ options ] <class> [ arguments ... ]
javaw [ options ] -jar <file.jar> [ arguments ... ]
```

Parameters

[options]

Command-line options to be passed to the runtime environment.

<class>

Startup class. The class must contain a `main()` method.

<file.jar>

Name of the `.jar` file to start. It is used only with the **-jar** option. The named `.jar` file must contain class and resource files for the application, with the startup class indicated by the `Main-Class` manifest header.

[arguments ...]

Command-line arguments to be passed to the `main()` function of the startup class.

Obtaining version information

You obtain the IBM build and version number for your Java installation using the **-version** or **-fullversion** options. You can also obtain version information for all jar files on the class path by using the **-Xjarversion** option.

Procedure

1. Open a shell prompt.
2. Type the following command:

```
java -version
```

You will see information similar to:

```
java version "1.6.0"  
Java(TM) SE Runtime Environment (build pmz6460sr10-20111027_02(SR10))  
IBM J9 VM (build 2.4, JRE 1.6.0 IBM J9 2.4 z/OS s390x-64 jvmmz6460sr10-20111026_93491 (JIT enabled,  
AOT enabled)  
J9VM - 20111026_093491  
JIT - r9_20111021_21134  
GC - 20110519_AA  
JCL - 20111025_01
```

Exact build dates and versions will change.

3. To obtain only the build information for the JVM, type the following command:

```
java -fullversion
```

You will see information similar to:

```
java full version "JRE 1.6.0 IBM Windows 32 build pwi3260sr11-20120412_01 (SR11)"
```

What to do next

You can also list the version information for all available jar files on the class path, the boot class path, and in the extensions directory. Type the following command:

```
java -Xjarversion
```

You will see information similar to:

```
...  
/usr/lpp/java/J6.0[_64]/jre/lib/ext/ibmpkcs11impl.jar VERSION: 1.0 build_20070125  
/usr/lpp/java/J6.0[_64]/jre/lib/ext/dfjview.jar  
/usr/lpp/java/J6.0[_64]/jre/lib/ext/xmlencfw.jar VERSION: 1.00, 20061011 LEVEL: -20061011  
...
```

The information available varies for each jar file and is taken from the **Implementation-Version** and **Build-Level** properties in the manifest of the jar file.

To query the Java version information programmatically, see “Determining which JVM version your application is running on” on page 38.

Specifying Java options and system properties

You can specify Java options and system properties directly on the command line. You can also use an options file or an environment variable.

About this task

The sequence of the Java options on the command line defines which options take precedence during startup. Rightmost options have precedence over leftmost options. In the following example, the **-Xjit** option takes precedence:

```
java -Xint -Xjit myClass
```

Use one or more of the options that are shown in the procedure to customize your runtime environment.

Procedure

1. Specify options or system properties on the command line. For example:

```
java -Dmysysprop1=tcip -Dmysysprop2=wait -Xdisablejavadump MyJavaClass
```
2. Create an environment variable that is called **IBM_JAVA_OPTIONS** containing the options. For example:

```
export IBM_JAVA_OPTIONS="-Dmysysprop1=tcip -Dmysysprop2=wait  
-Xdisablejavadump"
```
3. Create a file that contains the options, and specify that file on the command line or in the **IBM_JAVA_OPTIONS** environment variable by using the **-Xoptionsfile** parameter. For more information about constructing this file, see “-Xoptionsfile” on page 89.

Standard options

The definitions for the standard options.

See “JVM command-line options” on page 82 for information about nonstandard (-X) options.

-agentlib:*<libname>*[=*<options>*]

Loads a native agent library *<libname>*; for example **-agentlib:hprof**. For more information, specify **-agentlib:jwp=help** and **-agentlib:hprof=help** on the command line.

-agentpath:*libname*[=*<options>*]

Loads a native agent library by full path name.

-cp *<directories and .zip or .jar files separated by :>*

Sets the search path for application classes and resources. If **-classpath** and **-cp** are not used and the **CLASSPATH** environment variable is not set, the user class path is, by default, the current directory (.).

-classpath *<directories and .zip or .jar files separated by :>*

Sets the search path for application classes and resources. If **-classpath** and **-cp** are not used and the **CLASSPATH** environment variable is not set, the user class path is, by default, the current directory (.).

-D*<property name>*==*<value>*

Sets a system property.

-help or **-?**

Prints a usage message.

-javaagent:*<jarpath>*[=*<options>*]

Load a Java programming language agent. For more information, see the `java.lang.instrument` API documentation.

-jre-restrict-search

Include user private JREs in the version search.

-no-jre-restrict-search

Exclude user private JREs in the version search.

-showversion

Prints product version and continues.

-verbose:*<option>*[,*<option>*...]

Enables verbose output. Separate multiple options using commas. The available options are:

class

Writes an entry to stderr for each class that is loaded.

gc Writes verbose garbage collection information to stderr. Use **-Xverbosegclog** (see “Garbage Collector command-line options” on page 101 for more information) to control the output. See Verbose garbage collection logging for more information.

jni

Writes information to stderr describing the JNI services called by the application and JVM.

sizes

Writes information to stderr describing the active memory usage settings.

stack

Writes information to stderr describing the Java and C stack usage for each thread.

-version

Prints product version.

-version:<value>

Requires the specified version to run, for example “1.5”.

-X Prints help on nonstandard options.

Globalization of the java command

The **java** and **javaw** launchers accept arguments and class names containing any character that is in the character set of the current locale. You can also specify any Unicode character in the class name and arguments by using Java escape sequences.

To do this, use the **-Xargencoding** command-line option.

-Xargencoding

Use argument encoding. To specify a Unicode character, use escape sequences in the form `\u####`, where # is a hexadecimal digit (0 to 9, A to F).

-Xargencoding:utf8

Use UTF8 encoding.

-Xargencoding:latin

Use ISO8859_1 encoding.

For example, to specify a class called HelloWorld using Unicode encoding for both capital letters, use this command:

```
java -Xargencoding '\u0048ello\u0057orld'
```

The **java** and **javaw** commands provide translated messages. These messages differ based on the locale in which Java is running. The detailed error descriptions and other debug information that is returned by **java** is in English.

The Just-In-Time (JIT) compiler

The IBM Just-In-Time (JIT) compiler dynamically generates machine code for frequently used bytecode sequences in Java applications and applets during their execution. The JIT compiler delivers new optimizations as a result of compiler research, improves optimizations implemented in previous versions of the JIT, and provides better hardware exploitation.

The JIT is included in both the IBM SDK and Runtime Environment, which is enabled by default in user applications and SDK tools. Typically, you do not start

the JIT explicitly; the compilation of Java bytecode to machine code occurs transparently. You can disable the JIT to help isolate a problem. If a problem occurs when executing a Java application or an applet, you can disable the JIT to help isolate the problem. Disabling the JIT is a temporary measure only; the JIT is required to optimize performance.

For more information about the JIT, see JIT and AOT problem determination.

Disabling the JIT

The JIT can be disabled in a number of different ways. Both command-line options override the **JAVA_COMPILER** environment variable.

About this task

Turning off the JIT is a temporary measure that can help isolate problems when debugging Java applications.

Procedure

- Set the **JAVA_COMPILER** environment variable to **NONE** or the empty string before running the **java** application. Type the following command at a shell prompt:

```
export JAVA_COMPILER=NONE
```
- Use the **-D** option on the JVM command line to set the **java.compiler** property to **NONE** or the empty string. Type the following command at a shell prompt:

```
java -Djava.compiler=NONE <class>
```
- Use the **-Xint** option on the JVM command line. Type the following command at a shell prompt:

```
java -Xint <class>
```

Enabling the JIT

The JIT is enabled by default. You can explicitly enable the JIT in a number of different ways. Both command-line options override the **JAVA_COMPILER** environment variable.

Procedure

- Set the **JAVA_COMPILER** environment variable to **jitc** before running the Java application. At a shell prompt, enter:

```
export JAVA_COMPILER=jitc
```

If the **JAVA_COMPILER** environment variable is an empty string, the JIT remains disabled. To disable the environment variable, at the prompt, enter:

```
unset JAVA_COMPILER
```

- Use the **-D** option on the JVM command line to set the **java.compiler** property to **jitc**. At a prompt, enter:

```
java -Djava.compiler=jitc <class>
```
- Use the **-Xjit** option on the JVM command line. Do **not** specify the **-Xint** option at the same time. At a prompt, enter:

```
java -Xjit <class>
```

Determining whether the JIT is enabled

You can determine the status of the JIT using the **-version** option.

Procedure

Run the **java** launcher with the **-version** option. Enter the following command at a shell prompt:

```
java -version
```

If the JIT is not in use, a message is displayed that includes the following text:
(JIT disabled)

If the JIT is in use, a message is displayed that includes the following text:
(JIT enabled)

What to do next

For more information about the JIT, see The JIT compiler.

Specifying a garbage collection policy

The Garbage Collector manages the memory used by Java and by applications running in the JVM.

When the Garbage Collector receives a request for storage, unused memory in the heap is set aside in a process called "allocation". The Garbage Collector also checks for areas of memory that are no longer referenced, and releases them for reuse. This is known as "collection".

The collection phase can be triggered by a memory allocation fault, which occurs when no space remains for a storage request, or by an explicit `System.gc()` call.

Garbage collection can significantly affect application performance, so the IBM virtual machine provides various methods of optimizing the way garbage collection is carried out, potentially reducing the effect on your application.

For more detailed information about garbage collection, see Detailed description of garbage collection.

Garbage collection options

The **-Xgcpolicy** options control the behavior of the Garbage Collector. They make trade-offs between throughput of the application and overall system, and the pause times that are caused by garbage collection.

The format of the option is as follows:

```
-Xgcpolicy:<value>
```

The following values are available:

gencon

The **gencon** policy uses a concurrent mark phase combined with generational garbage collection to help minimize the time that is spent in any garbage collection pause. This policy is particularly useful for applications with many short-lived objects, such as transactional applications. Pause times can be significantly shorter than with the **optthroughput** policy, while still producing good throughput. Heap fragmentation is also reduced.

optavgpause

The **optavgpause** policy uses concurrent mark and concurrent sweep phases.

Pause times are shorter than with **optthruput**, but application throughput is reduced because some garbage collection work is taking place while the application is running. Consider using this policy if you have a large heap size (available on 64-bit platforms), because this policy limits the effect of increasing heap size on the length of the garbage collection pause. However, if your application uses many short-lived objects, the **gencon** policy might produce better performance.

subpool

The **subpool** policy disables the concurrent mark phase, but employs an improved object allocation algorithm to achieve better performance when allocating objects on the heap. This algorithm is more suitable for multiple processor systems, commonly 16 processors or more. Applications that must scale on large systems might benefit from this policy. This policy is available on AIX®, Linux PPC and zSeries, z/OS, and i5/OS™ only.

optthruput

The **optthruput** policy (default) disables the concurrent mark phase. The application stops during global garbage collection, so long pauses can occur. This configuration is typically used for large-heap applications when high application throughput, rather than short garbage collection pauses, is the main performance goal. If your application cannot tolerate long garbage collection pauses, consider using another policy, such as **gencon**.

More effective heap usage using compressed references

Many Java application workloads depend on the Java heap size. The IBM SDK for Java can use compressed references on 64-bit platforms to decrease the size of Java objects and make more effective use of the available space. The result is less frequent garbage collection and improved memory cache utilization.

If you specify the **-Xnocompressedrefs** command-line option, the IBM SDK for Java 64-bit stores object references as 64-bit values. If you specify the **-Xcompressedrefs** command-line option, object references are stored as 32-bit representation, which reduces the 64-bit object size to be the same as a 32-bit object.

As the 64-bit objects with compressed references are smaller than default 64-bit objects, they occupy a smaller memory footprint in the Java heap and improves data locality. This results in better memory utilization and improved performance.

If you are using a 64-bit IBM SDK for Java, use **-Xcompressedrefs** whenever you require a maximum heap size of less than 25 GB. For example, your application might use a lot of native memory and require the JVM to run in a small footprint.

Note: If you are using compressed references on z/OS v1.10 or earlier, you must use APAR OA26294.

See Compressed references for more detailed information and hardware/operating system specific guidance on compressed references. More information is also available in the Websphere white paper on compressed references.

Pause time

If an object cannot be created from the available space in the heap, the Garbage Collector attempts to tidy the heap. The intention is that subsequent allocation requests can be satisfied quickly.

The Garbage Collector tries to returning the heap to a state in which the immediate and subsequent space requests are successful. The Garbage Collector identifies unreferenced “garbage” objects, and deletes them. This work takes place in a garbage collection cycle. These cycles might introduce occasional, unexpected pauses in the execution of application code. As applications grow in size and complexity, and heaps become correspondingly larger, the garbage collection pause time tends to grow in size and significance. Pause time can vary from a few milliseconds to many seconds. The actual time depends on the size of the heap, and the quantity of garbage.

Pause time reduction

The JVM uses two techniques to reduce pause times: concurrent garbage collection and generational garbage collection.

The **-Xgcpolicy:optavgpause** command-line option requests the use of concurrent garbage collection (GC) to reduce significantly the time that is spent in garbage collection pauses. Concurrent GC reduces the pause time by performing some garbage collection activities concurrently with normal program execution to minimize the disruption caused by the collection of the heap. The **-Xgcpolicy:optavgpause** option also limits the effect of increasing the heap size on the length of the garbage collection pause. The **-Xgcpolicy:optavgpause** option is most useful for configurations that have large heaps. With the reduced pause time, you might experience some reduction of throughput to your applications.

During concurrent GC, a significant amount of time is wasted identifying relatively long-lasting objects that cannot then be collected. If garbage collection concentrates on only the objects that are most likely to be recyclable, you can further reduce pause times for some applications. Generational GC reduces pause times by dividing the heap into two generations: the “new” and the “tenure” areas. Objects are placed in one of these areas depending on their age. The new area is the smaller of the two and contains new objects; the tenure is larger and contains older objects. Objects are first allocated to the new area; if they have active references for long enough, they are promoted to the tenure area.

Generational GC depends on most objects not lasting long. Generational GC reduces pause times by concentrating the effort to reclaim storage on the new area because it has the most recyclable space. Rather than occasional but lengthy pause times to collect the entire heap, the new area is collected more frequently and, if the new area is small enough, pause times are comparatively short. However, generational GC has the drawback that, over time, the tenure area might become full. To minimize the pause time when this situation occurs, use a combination of concurrent GC and generational GC. The **-Xgcpolicy:gencon** option requests the combined use of concurrent and generational GC to help minimize the time that is spent in any garbage collection pause.

Environments with very full heaps

If the Java heap becomes nearly full, and very little garbage can be reclaimed, requests for new objects might not be satisfied quickly because no space is immediately available.

If the heap is operated at near-full capacity, application performance might suffer regardless of which garbage collection options are used; and, if requests for more heap space continue to be made, the application might receive an `OutOfMemoryError`, which results in JVM termination if the exception is not caught and handled. At this point, the JVM produces a Javadump file for use

during diagnostic procedures. In these conditions, you are recommended either to increase the heap size by using the `-Xmx` option or to reduce the number of objects in use.

For more information, see Garbage Collector diagnostic data.

Euro symbol support

The IBM SDK and Runtime Environment set the Euro as the default currency for those countries in the European Monetary Union (EMU) for dates on or after 1 January, 2002. From 1 January 2008, Cyprus and Malta also have the Euro as the default currency.

To use the old national currency, specify `-Duser.variant=PREEURO` on the Java command line.

If you are running the UK, Danish, or Swedish locales and want to use the Euro, specify `-Duser.variant=EURO` on the Java command line.

Using Indian and Thai input methods

From Version 6, the Indian and Thai input methods are not available by default. You must manually include the input method jar files in your Java extensions path to use the Indian and Thai input methods.

About this task

In Version 5.0, the input method jar files were included in the `jre/lib/ext` directory and were automatically loaded by the JVM. In Version 6, the input method jar files are included in the `jre/lib/im` directory and you must manually add them to the Java extensions path to enable Indian and Thai input methods.

Procedure

- Copy the `indicim.jar` and `thaiim.jar` files from the `jre/lib/im` directory to the `jre/lib/ext` directory.
- Add the `jre/lib/im` directory to the extension directories system property. Use the following command-line option:

```
java -Djava.ext.dirs=/usr/lpp/java/J6.0[_64]/jre/lib/ext:  
/usr/lpp/java/J6.0[_64]/jre/lib/im <class>
```

What to do next

If you installed the SDK or Runtime Environment in a different directory, replace `/usr/lpp/java/J6.0[_64]/` with the directory in which you installed the SDK or Runtime Environment.

Chapter 5. Developing Java applications

The SDK contains many tools and libraries required for Java software development.

See “Contents of the SDK” on page 6 for details of the tools available.

Using XML

The IBM SDK contains the XML4J and XL XP-J parsers, the XL TXE-J 1.0 XSLT compiler, and the XSLT4J XSLT interpreter. These tools allow you to parse, validate, transform, and serialize XML documents independently from any given XML processing implementation.

Use factory finders to locate implementations of the abstract factory classes, as described in “Selecting an XML processor” on page 28. By using factory finders, you can select a different XML library without changing your Java code.

Available XML libraries

The IBM SDK for Java contains the following XML libraries:

XML4J 4.5

XML4J is a validating parser providing support for the following standards:

- XML 1.0 (4th edition)
- Namespaces in XML 1.0 (2nd edition)
- XML 1.1 (2nd edition)
- Namespaces in XML 1.1 (2nd edition)
- W3C XML Schema 1.0 (2nd Edition)
- XInclude 1.0 (2nd Edition)
- OASIS XML Catalogs 1.0
- SAX 2.0.2
- DOM Level 3 Core, Load and Save
- DOM Level 2 Core, Events, Traversal and Range
- JAXP 1.4

XML4J 4.5 is based on Apache Xerces-J 2.9.0. See <http://xerces.apache.org/xerces2-j/> for more information.

XL XP-J 1.1

XL XP-J 1.1 is a high-performance non-validating parser that provides support for StAX 1.0 (JSR 173). StAX is a bidirectional API for pull-parsing and streaming serialization of XML 1.0 and XML 1.1 documents. See the “XL XP-J reference information” on page 31 section for more details about what is supported by XL XP-J 1.1.

XL TXE-J 1.0

For Version 5.0, the IBM SDK for Java included the XSLT4J compiler and interpreter. The XSLT4J interpreter was used by default.

For Version 6 and later, the IBM SDK for Java includes XL TXE-J. XL TXE-J includes the XSLT4J 2.7.8 interpreter and a new XSLT compiler. The new compiler is used by default. The XSLT4J compiler is no longer included with the IBM SDK for Java. See “Migrating to the XL-TXE-J” on page 29 for information about migrating to XL TXE-J.

XL TXE-J provides support for the following standards:

- XSLT 1.0
- XPath 1.0
- JAXP 1.4

Selecting an XML processor

XML processor selection is performed using service providers. When using a factory finder, Java looks in the following places, in this order, to see which service provider to use:

1. The system property with the same name as the service provider.
2. The service provider specified in a properties file.
 - **For XMLEventFactory, XMLInputFactory, and XMLOutputFactory only.** The value of the service provider in the file `/usr/lpp/java/J6.0[_64]/jre/lib/stax.properties`.
 - **For other factories.** The value of the service provider in the file `/usr/lpp/java/J6.0[_64]/jre/lib/jaxp.properties`.
3. The contents of the `META-INF/services/<service.provider>` file.
4. The default service provider.

The following service providers control the XML processing libraries used by Java:

javax.xml.parsers.SAXParserFactory

Selects the SAX parser. By default, `org.apache.xerces.jaxp.SAXParserFactoryImpl` from the XML4J library is used.

javax.xml.parsers.DocumentBuilderFactory

Selects the document builder. By default, `org.apache.xerces.jaxp.DocumentBuilderFactoryImpl` from the XML4J library is used.

javax.xml.datatype.DatatypeFactory

Selects the datatype factory. By default, `org.apache.xerces.jaxp.datatype.DatatypeFactoryImpl` from the XML4J library is used.

javax.xml.stream.XMLEventFactory

Selects the StAX event factory. By default, `com.ibm.xml.xlcp.api.stax.XMLEventFactoryImpl` from the XL XP-J library is used.

javax.xml.stream.XMLInputFactory

Selects the StAX parser. By default, `com.ibm.xml.xlcp.api.stax.XMLInputFactoryImpl` from the XL XP-J library is used.

javax.xml.stream.XMLOutputFactory

Selects the StAX serializer. By default, `com.ibm.xml.xlcp.api.stax.XMLOutputFactoryImpl` from the XL XP-J library is used.

javax.xml.transform.TransformerFactory

Selects the XSLT processor. Possible values are:

com.ibm.xtq.xslt.jaxp.compiler.TransformerFactoryImpl

Use the XL TXE-J compiler. This value is the default.

org.apache.xalan.processor.TransformerFactoryImpl

Use the XSLT4J interpreter.

javax.xml.validation.SchemaFactory:http://www.w3.org/2001/XMLSchema

Selects the schema factory for the W3C XML Schema language. By default, org.apache.xerces.jaxp.validation.XMLSchemaFactory from the XML4J library is used.

javax.xml.xpath.XPathFactory

Selects the XPath processor. By default, org.apache.xpath.jaxp.XPathFactoryImpl from the XSLT4J library is used.

Migrating to the XL-TXE-J

From Version 6, the XL TXE-J compiler replaces the XSLT4J interpreter as the default XSLT processor. If you are migrating applications from older versions of Java, follow these steps to prepare your application for the new library.

About this task

The XL TXE-J compiler is faster than the XSLT4J interpreter when you are applying the same transformation more than once. If you perform each individual transformation only once, the XL TXE-J compiler is slower than the XSLT4J interpreter because compilation and optimization reduce performance.

To continue using the XSLT4J interpreter as your XSLT processor, set the **javax.xml.transform.TransformerFactory** service provider to org.apache.xalan.processor.TransformerFactoryImpl.

To migrate to the XL-TXE-J compiler, follow the instructions in this task.

Procedure

1. Use com.ibm.xtq.xslt.jaxp.compiler.TransformerFactoryImpl when setting the **javax.xml.transform.TransformerFactory** service provider.
2. Regenerate class files generated by the XSLT4J compiler. XL TXE-J cannot execute class files generated by the XSLT4J compiler.
3. Some methods generated by the compiler might exceed the JVM method size limit, in which case the compiler attempts to split these methods into smaller methods.
 - If the compiler splits the method successfully, you receive the following warning:
Some generated functions exceeded the JVM method size limit and were automatically split into smaller functions. You might get better performance by manually splitting very large templates into smaller templates, by using the 'splitlimit' option to the Process or Compile command, or by setting the 'http://www.ibm.com/xmlns/prod/xltxe-j/split-limit' transformer factory attribute. You can use the compiled classes, but you might get better performance by controlling the split limit manually.
 - If the compiler does not split the method successfully, you receive one of the following exceptions:

com.ibm.xtq.bcel.generic.ClassGenException: Branch target offset too large for short or
bytecode array size > 65535 at offset=#####Try setting the split limit manually, or decreasing the split limit.

To set the split limit, use the **-SPLITLIMIT** option when using the **Process** or **Compile** commands, or the <http://www.ibm.com/xmlns/prod/xltxe-j/split-limit> transformer factory attribute when using the transformer factory. The split limit can be between 100 and 2000. When setting the split limit manually, use the highest split limit possible for best performance.

4. XL TXE-J might need more memory than the XSLT4J compiler. If you are running out of memory or performance seems slow, increase the size of the heap using the **-Xmx** option.
5. Migrate your application to use the new attribute keys. The old transformer factory attribute keys are deprecated. The old names are accepted with a warning.

Table 3. Changes to attribute keys from the XSL4J compiler to the XL TXE-J compiler

XSL4J compiler attribute	XL TXE-J compiler attribute
translet-name	http://www.ibm.com/xmlns/prod/xltxe-j/translet-name
destination-directory	http://www.ibm.com/xmlns/prod/xltxe-j/destination-directory
package-name	http://www.ibm.com/xmlns/prod/xltxe-j/package-name
jar-name	http://www.ibm.com/xmlns/prod/xltxe-j/jar-name
generate-translet	http://www.ibm.com/xmlns/prod/xltxe-j/generate-translet
auto-translet	http://www.ibm.com/xmlns/prod/xltxe-j/auto-translet
use-classpath	http://www.ibm.com/xmlns/prod/xltxe-j/use-classpath
debug	http://www.ibm.com/xmlns/prod/xltxe-j/debug
indent-number	http://www.ibm.com/xmlns/prod/xltxe-j/indent-number
enable-inlining	<i>Obsolete in new compiler</i>

6. Optional: For best performance, ensure that you are not recompiling XSLT transformations that can be reused. Use one of the following methods to reuse compiled transformations:
 - If your stylesheet does not change at run time, compile the stylesheet as part of your build process and put the compiled classes on your classpath. Use the `org.apache.xalan.xsltc.cmdline.Compile` command to compile the stylesheet and set the <http://www.ibm.com/xmlns/prod/xltxe-j/use-classpath> transformer factory attribute to true to load the classes from the classpath.
 - If your application will use the same stylesheet during multiple runs, set the <http://www.ibm.com/xmlns/prod/xltxe-j/auto-translet> transformer factory attribute to true to automatically save the compiled stylesheet to disk for reuse. The compiler will use a compiled stylesheet if it is available, and compile the stylesheet if it is not available or is out-of-date. Use the <http://www.ibm.com/xmlns/prod/xltxe-j/destination-directory> transformer factory attribute to set the directory used to store compiled stylesheets. By default, compiled stylesheets are stored in the same directory as the stylesheet.
 - If your application is a long-running application that reuses the same stylesheet, use the transformer factory to compile the stylesheet and create a

Templates object. You can use the Templates object to create Transformer objects without recompiling the stylesheet. The Transformer objects can also be reused but are not thread-safe.

- If your application uses each stylesheet just once or a very small number of times, or you are unable to make any of the other changes listed in this step, you might want to continue to use the XSLT4J interpreter by setting the `javax.xml.transform.TransformerFactory` service provider to `org.apache.xalan.processor.TransformerFactoryImpl`.

Securing JAXP processing against malformed input

If your application takes untrusted XML, XSD or XSL files as input, you can enforce specific limits during JAXP processing to protect your application from malformed data. If you specify limits, you must override the default XML parser configuration with a custom configuration.

About this task

To protect your application from malformed data, you can enforce specific limits during JAXP processing. These limits can be set in your `jaxp.properties` file, or by specifying various system properties on the command line. However, for these limits to take effect you must also override the default XML parser configuration with a custom configuration that allows these secure processing limits.

Procedure

1. Select the limits that you want to set for your application.
 - a. To limit the number of entity expansions in an XML document, see “-Djdk.xml.entityExpansionLimit” on page 78.
 - b. To limit the maximum size of a general entity, see “-Djdk.xml.maxGeneralEntitySizeLimit” on page 79.
 - c. To limit the maximum size of a parameter entity, see “-Djdk.xml.maxParameterEntitySizeLimit” on page 80
 - d. To limit the total size of all entities that include general and parameter entities, see “-Djdk.xml.totalEntitySizeLimit” on page 81
 - e. To define the maximum number of content model nodes that can be created in a grammar, see “-Djdk.xml.maxOccur” on page 79
2. To override the default XML parser configuration, set the custom configuration by specifying the following system property on the command line:
-Dorg.apache.xerces.xni.parser.XMLParserConfiguration=*config_file*, where *config_file* is **org.apache.xerces.parsers.SecureProcessingConfiguration**. For more information about the full override mechanism, see <http://xerces.apache.org/xerces2-j/faq-xni.html#faq-2>.

XML reference information

The XL XP-J and XL TXE-J XML libraries are new for Version 6 of the SDK. This reference information describes the features supported by these libraries.

XL XP-J reference information

XL XP-J 1.1 is a high-performance non-validating parser that provides support for StAX 1.0 (JSR 173). StAX is a bidirectional API for pull-parsing and streaming serialization of XML 1.0 and XML 1.1 documents.

Unsupported features

The following optional StAX features are not supported by XL XP-J:

- DTD validation when using an XMLStreamReader or XMLEventReader. The XL XP-J parser is non-validating.
- When using an XMLStreamReader to read from a character stream (java.io.Reader), the Location.getCharaterOffset() method always returns -1. The Location.getCharaterOffset() returns the byte offset of a Location when using an XMLStreamReader to read from a byte stream (java.io.InputStream).

XMLInputFactory reference

The javax.xml.stream.XMLInputFactory implementation supports the following properties, as described in the XMLInputFactory Javadoc information:

<http://download.oracle.com/javase/6/docs/api/javax/xml/stream/XMLInputFactory.html>.

Property name	Supported?
<code>javax.xml.stream.isValidating</code>	No. The XL XP-J scanner does not support validation.
<code>javax.xml.stream.isNamespaceAware</code>	Yes, supports true and false. For XMLStreamReaders created from DOMSources, namespace processing depends on the methods that were used to create the DOM tree, and this value has no effect.
<code>javax.xml.stream.isCoalescing</code>	Yes
<code>javax.xml.stream.isReplacingEntityReferences</code>	Yes. For XMLStreamReaders created from DOMSources, if entities have already been replaced in the DOM tree, setting this parameter has no effect.
<code>javax.xml.stream.isSupportingExternalEntities</code>	Yes
<code>javax.xml.stream.supportDTD</code>	<p>True is always supported. Setting the value to false works only if the <code>com.ibm.xml.xlxp.support.dtd.compat.mode</code> system property is also set to false.</p> <p>When both properties are set to false, parsers created by the factory throw an XMLStreamException when they encounter an entity reference that requires expansion. This setting is useful for protecting against Denial of Service (DoS) attacks involving entities declared in the DTD.</p> <p>Setting the value to false does not work before Service Refresh 2.</p>
<code>javax.xml.stream.reporter</code>	Yes
<code>javax.xml.stream.resolver</code>	Yes

XL XP-J also supports the optional method `createXMLStreamReader(javax.xml.transform.Source)`, which allows StAX readers to be created from DOM and SAX sources.

XL XP-J also supports the `javax.xml.stream.isSupportingLocationCoordinates` property. If you set this property to true, XMLStreamReaders created by the factory return accurate line, column, and character information using Location objects. If you set this property to false, line, column, and character information is not available. By default, this property is set to false for performance reasons.

XMLStreamReader reference

The `javax.xml.stream.XMLStreamReader` implementation supports the following properties, as described in the XMLStreamReader Javadoc: <http://download.oracle.com/javase/6/docs/api/javax/xml/stream/XMLStreamReader.html>.

Property name	Supported?
<code>javax.xml.stream.entities</code>	Yes
<code>javax.xml.stream.notations</code>	Yes

XL XP-J also supports the `javax.xml.stream.isInterning` property. This property returns a boolean value indicating whether or not XML names and namespace URIs returned by the API calls have been interned by the parser. This property is read-only.

XMLOutputFactory reference

The `javax.xml.stream.XMLOutputFactory` implementation supports the following properties, as described in the XMLOutputFactory Javadoc: <http://download.oracle.com/javase/6/docs/api/javax/xml/stream/XMLOutputFactory.html>.

Property name	Supported?
<code>javax.xml.stream.isRepairingNamespaces</code>	Yes

XL XP-J also supports the `javax.xml.stream.XMLOutputFactory.recycleWritersOnEndDocument` property. If you set this property to true, XMLStreamWriters created by this factory are recycled when `writeEndDocument()` is called. After recycling, some XMLStreamWriter methods, such as `getNamespaceContext()`, must not be called. By default, XMLStreamWriters are recycled when `close()` is called. You must call the `XMLStreamWriter.close()` method when you have finished with an XMLStreamWriter, even if this property is set to true.

XMLStreamWriter reference

The `javax.xml.stream.XMLStreamWriter` implementation supports the following properties, as described in the XMLStreamWriter Javadoc: <http://download.oracle.com/javase/6/docs/api/javax/xml/stream/XMLStreamWriter.html>.

Property name	Supported?
<code>javax.xml.stream.isRepairingNamespaces</code>	Yes

Properties on XMLStreamWriter objects are read-only.

XL XP-J also supports the `javax.xml.stream.XMLStreamWriter.isSetPrefixBeforeStartElement` property. This property returns a Boolean indicating whether calls to `setPrefix()` and `setDefaultNamespace()` should occur before calls to `writeStartElement()` or `writeEmptyElement()` to put a namespace prefix in scope for that element. XL XP-J always returns false; calls to `setPrefix()` and `setDefaultNamespace()` should occur

after writeStartElement() or writeEmptyElement().

XL TXE-J reference information

XL TXE-J is an XSLT library containing the XSLT4J 2.7.8 interpreter and a XSLT compiler.

Feature comparison table

Table 4. Comparison of the features in the XSLT4J interpreter, the XSLT4J compiler, and the XL TXE-J compiler.

Feature	XSLT4J interpreter (included)	XSLT4J compiler (not included)	XL TXE-J compiler (included)
http://javax.xml.transform.stream.StreamSource/feature feature	Yes	Yes	Yes
http://javax.xml.transform.stream.StreamResult/feature feature	Yes	Yes	Yes
http://javax.xml.transform.dom.DOMSource/feature feature	Yes	Yes	Yes
http://javax.xml.transform.dom.DOMResult/feature feature	Yes	Yes	Yes
http://javax.xml.transform.sax.SAXSource/feature feature	Yes	Yes	Yes
http://javax.xml.transform.sax.SAXResult/feature feature	Yes	Yes	Yes
http://javax.xml.transform.stax.StAXSource/feature feature	Yes	No	Yes
http://javax.xml.transform.stax.StAXResult/feature feature	Yes	No	Yes
http://javax.xml.transform.sax.SAXTransformerFactory/feature feature	Yes	Yes	Yes
http://javax.xml.transform.sax.SAXTransformerFactory/feature/xmlfilter feature	Yes	Yes	Yes
http://javax.xml.XMLConstants/feature/secure-processing feature	Yes	Yes	Yes
http://xml.apache.org/xalan/features/incremental attribute	Yes	No	No
http://xml.apache.org/xalan/features/optimize attribute	Yes	No	No
http://xml.apache.org/xalan/properties/source-location attribute	Yes	No	No
translet-name attribute	N/A	Yes	Yes (with new name)
destination-directory attribute	N/A	Yes	Yes (with new name)
package-name attribute	N/A	Yes	Yes (with new name)
jar-name attribute	N/A	Yes	Yes (with new name)
generate-translet attribute	N/A	Yes	Yes (with new name)

Table 4. Comparison of the features in the XSLT4J interpreter, the XSLT4J compiler, and the XL TXE-J compiler. (continued)

Feature	XSLT4J interpreter (included)	XSLT4J compiler (not included)	XL TXE-J compiler (included)
auto-translet attribute	N/A	Yes	Yes (with new name)
use-classpath attribute	N/A	Yes	Yes (with new name)
enable-inlining attribute	No	Yes	No (obsolete in TL TXE-J)
indent-number attribute	No	Yes	Yes (with new name)
debug attribute	No	Yes	Yes (with new name)
Java extensions	Yes	Yes (abbreviated syntax only, xalan:component/xalan:script constructs not supported)	Yes (abbreviated syntax only, xalan:component/xalan:script constructs not supported)
JavaScript extensions	Yes	No	No
Extension elements	Yes	No	No
EXSLT extension functions	Yes	Yes (excluding dynamic)	Yes (excluding dynamic)
redirect extension	Yes	Yes (excluding redirect:open and redirect:close)	Yes
output extension	No	Yes	Yes
nodeset extension	Yes	Yes	Yes
NodeInfo extension functions	Yes	No	No
SQL library extension	Yes	No	No
pipeDocument extension	Yes	No	No
evaluate extension	Yes	No	No
tokenize extension	Yes	No	No
XML 1.1	Yes	Yes	Yes

Notes

1. With the Process command, use **-FLAVOR sr2sw** to transform using StAX stream processing, and **-FLAVOR er2ew** for StAX event processing.
2. The new compiler does not look for the org.apache.xalan.xsltc.dom.XSLTCDTMMManager service provider. Instead, if StreamSource is used, the compiler switches to a high-performance XML parser.
3. Inlining is obsolete in XL TXE-J.
 - The **-XN** option to the **Process** command is silently ignored.
 - The **-n** option to the **Compile** command is silently ignored.
 - The **enable-inlining** transformer factory attribute is silently ignored.
4. The org.apache.xalan.xsltc.trax.SmartTransformerFactoryImpl class is no longer supported.

Using an older version of Xerces or Xalan

If you are using an older version of Xerces (before 2.0) or Xalan (before 2.3) in the endorsed override, you might get a `NullPointerException` when you start your application. This exception occurs because these older versions do not handle the `jaxp.properties` file correctly.

About this task

To avoid this situation, use one of the following workarounds:

- Upgrade to a newer version of the application that implements the latest Java API for XML Programming (JAXP) specification (<https://jaxp.dev.java.net/>).
- Remove the `jaxp.properties` file from `/usr/lib/java/J6.0[_64]/jre/lib`.
- Uncomment the entries in the `jaxp.properties` file in `/usr/lib/java/J6.0[_64]/jre/lib`.
- Set the system property for `javax.xml.parsers.SAXParserFactory`, `javax.xml.parsers.DocumentBuilderFactory`, or `javax.xml.transform.TransformerFactory` using the `-D` command-line option.
- Set the system property for `javax.xml.parsers.SAXParserFactory`, `javax.xml.parsers.DocumentBuilderFactory`, or `javax.xml.transform.TransformerFactory` in your application. For an example, see the JAXP 1.4 specification.
- Explicitly set the SAX parser, Document builder, or Transformer factory using the `IBM_JAVA_OPTIONS` environment variable.

```
export IBM_JAVA_OPTIONS=-Djavax.xml.parsers.SAXParserFactory=
org.apache.xerces.jaxp.SAXParserFactoryImpl
```

or

```
export IBM_JAVA_OPTIONS=-Djavax.xml.parsers.DocumentBuilderFactory=
org.apache.xerces.jaxp.DocumentBuilderFactoryImpl
```

or

```
export IBM_JAVA_OPTIONS=-Djavax.xml.transform.TransformerFactory=
org.apache.xalan.processor.TransformerFactoryImpl
```

Debugging Java applications

To debug Java programs, you can use the Java Debugger (JDB) application or other debuggers that communicate by using the Java Platform Debugger Architecture (JPDA) that is provided by the SDK for the operating system.

More information about problem diagnosis using Java can be found in the Diagnostic Guide.

Java Debugger (JDB)

The Java Debugger (JDB) is included in the SDK. The debugger is started with the `jdb` command; it attaches to the JVM using JPDA.

To debug a Java application:

1. Start the JVM with the following options:

```
java -agentlib:jdpw=transport=dt_socket,server=y,address=<port> <class>
```

The JVM starts up, but suspends execution before it starts the Java application.

2. In a separate session, you can attach the debugger to the JVM:

```
jdb -attach <port>
```

The debugger will attach to the JVM, and you can now issue a range of commands to examine and control the Java application; for example, type `run` to allow the Java application to start.

For more information about JDB options, type:

```
jdb -help
```

For more information about JDB commands:

1. Type `jdb`
2. At the **jdb** prompt, type `help`

You can also use JDB to debug Java applications running on remote workstations. JPDA uses a TCP/IP socket to connect to the remote JVM.

1. Start the JVM with the following options:

```
java -agentlib:jdwp=transport=dt_socket,server=y,address=<port> <class>
```

The JVM starts up, but suspends execution before it starts the Java application.

2. Attach the debugger to the remote JVM:

```
jdb -attach <host>:<port>
```

The Java Virtual Machine Debugging Interface (JVMDI) is not supported in this release. It has been replaced by the Java Virtual Machine Tool Interface (JVMTI).

For more information about JDB and JPDA and their usage, see these Web sites:

- <http://www.oracle.com/technetwork/java/javase/tech/jpda-141715.html>
- <http://download.oracle.com/javase/6/docs/technotes/guides/jpda/>
- <http://download.oracle.com/javase/6/docs/technotes/guides/jpda/jdb.html>

Selective debugging

Use the `com.ibm.jvm.Debuggable` annotation to mark classes and methods that should be available for debugging. Use the `-XselectiveDebug` parameter to enable selective debugging at run time. The JVM optimizes methods that do not need debugging to provide better performance in a debugging environment.

About this task

Selective debugging is useful when Java is being used as a framework for development, for example, as an IDE. The Java code for the IDE is optimized for performance while the user code is debugged.

Procedure

1. Import the `Debuggable` annotation from the `com.ibm.jvm` package.

```
import com.ibm.jvm.Debuggable;
```

2. Decorate methods using the `Debuggable` annotation.

```
@Debuggable
public int method1() {
    ...
}
```

3. Optional: You can also decorate classes using the `Debuggable` annotation. All methods in the class will remain debuggable.

```
@Debuggable
public class Class1 {
    ...
}
```

4. Enable selective debugging at run time using the `-XselectiveDebug` command-line option.

Results

Applications will run faster while being debugged because the core Java API and any IDE code can be optimized for performance.

Determining whether your application is running on a 31-bit or 64-bit JVM

Some Java applications must be able to determine whether they are running on a 31-bit JVM or on a 64-bit JVM. For example, if your application has a native code library, the library must be compiled separately in 31- and 64-bit forms for platforms that support both 31- and 64-bit modes of operation. In this case, your application must load the correct library at run environment time, because it is not possible to mix 31- and 64-bit code.

About this task

The system property `com.ibm.vm.bitmode` allows applications to determine the mode in which your JVM is running. It returns the following values:

- 32 - the JVM is running in 31-bit mode
- 64 - the JVM is running in 64-bit mode

You can inspect the `com.ibm.vm.bitmode` property from inside your application code using the call:

```
System.getProperty("com.ibm.vm.bitmode");
```

Determining which JVM version your application is running on

You can programmatically determine which JVM version your application is running on by querying the `java.runtime.version` system property.

About this task

The system property `java.runtime.version` can be queried at run time to determine the version of the JVM that is running.

Procedure

Use the following call: `System.getProperty("java.runtime.version");`

Results

This call returns a Java String with content similar to:

```
pwi3260sr11-20120412_01 (SR11)
```

The result of this query is similar to the command-line option `java -fullversion`. For more information about finding the Java version using command-line options, see “Obtaining version information” on page 17.

How the JVM processes signals

When a signal is raised that is of interest to the JVM, a signal handler is called. This signal handler determines whether it has been called for a Java or non-Java thread.

If the signal is for a Java thread, the JVM takes control of the signal handling. If an application handler for this signal is installed and you did not specify the **-Xnosigchain** command-line option, the application handler for this signal is called after the JVM has finished processing.

If the signal is for a non-Java thread, and the application that installed the JVM had previously installed its own handler for the signal, control is given to that handler. Otherwise, if the signal is requested by the JVM or Java application, the signal is ignored or the default action is taken.

For exception and error signals, the JVM either:

- Handles the condition and recovers, or
- Enters a controlled shut down sequence where it:
 1. Produces dumps, to describe the JVM state at the point of failure
 2. Calls your application's signal handler for that signal
 3. Calls any application-installed unexpected shut down hook
 4. Performs the necessary JVM cleanup

For information about writing a launcher that specifies the previous hooks, see: <http://www.ibm.com/developerworks/java/library/i-signalhandling/>. This item was written for Java V1.3.1, but still applies to later versions.

For interrupt signals, the JVM also enters a controlled shut down sequence, but this time it is treated as a normal termination that:

1. Calls your application's signal handler for that signal
2. Calls all application shut down hooks
3. Calls any application-installed exit hook
4. Performs the necessary JVM cleanup

The shut down is identical to the shut down initiated by a call to the Java method `System.exit()`.

Other signals that are used by the JVM are for internal control purposes and do not cause it to stop. The only control signal of interest is SIGQUIT, which causes a Javadump to be generated.

Signals used by the JVM

The types of signals are Exceptions, Errors, Interrupts, and Controls.

Table 5 on page 40 shows the signals that are used by the JVM. The signals are grouped in the table by type or use, as follows:

Exceptions

The operating system synchronously raises an appropriate exception signal whenever an unrecoverable condition occurs.

Errors The JVM raises a SIGABRT if it detects a condition from which it cannot recover.

Interrupts

Interrupt signals are raised asynchronously, from outside a JVM process, to request shut down.

Controls

Other signals that are used by the JVM for control purposes.

Table 5. Signals used by the JVM

Signal Name	Signal type	Description	Disabled by -Xrs	Disabled by -Xrs:sync
SIGBUS (7)	Exception	Incorrect access to memory (data misalignment)	Yes	Yes
SIGSEGV (11)	Exception	Incorrect access to memory (write to inaccessible memory)	Yes	Yes
SIGILL (4)	Exception	Illegal instruction (attempt to call an unknown machine instruction)	Yes	Yes
SIGFPE (8)	Exception	Floating point exception (divide by zero)	Yes	Yes
SIGABRT (6)	Error	Abnormal termination. The JVM raises this signal whenever it detects a JVM fault.	Yes	Yes
SIGINT (2)	Interrupt	Interactive attention (CTRL-C). JVM exits normally.	Yes	No
SIGTERM (15)	Interrupt	Termination request. JVM will exit normally.	Yes	No
SIGHUP (1)	Interrupt	Hang up. JVM exits normally.	Yes	No
SIGQUIT (3)	Control	By default, this triggers a Javadump.	Yes	No
SIGRECONFIG (58)	Control	Reserved to detect any change in the number of CPUs, processing capacity, or physical memory.	Yes	No
SIGTRAP (5)	Control	Used by the JIT.	Yes	Yes

Table 5. Signals used by the JVM (continued)

Signal Name	Signal type	Description	Disabled by -Xrs	Disabled by -Xrs:sync
SIGCHLD (17)	Control	Used by the SDK for internal control.	No	No
SIGUSR1	Control	Used by the SDK.	No	No

Note: A number supplied after the signal name is the standard numeric value for that signal.

JVM signal handlers for SIGSEGV/SIGILL/SIGBUS/SIGFPE/SIGTRAP and for SIGABRT by default terminate the process using `exit()`. If you are using the IBM z/OS Language Environment (LE), LE is not aware that the JVM ended abnormally. Use the `-Xsignal:posixSignalHandler=cooperativeShutdown` option to control how the signal handlers end. See “JVM command-line options” on page 82.

Use the `-Xrs` (reduce signal usage) option to prevent the JVM from handling most signals. For more information, see Oracle’s Java application launcher page.

Signals 1 (SIGHUP), 2 (SIGINT), 4 (SIGILL), 7 (SIGBUS), 8 (SIGFPE), 11 (SIGSEGV), and 15 (SIGTERM) on JVM threads cause the JVM to shut down; therefore, an application signal handler should not attempt to recover from these unless it no longer requires the JVM.

Linking a native code driver to the signal-chaining library

The Runtime Environment contains signal-chaining. Signal-chaining enables the JVM to interoperate more efficiently with native code that installs its own signal handlers.

About this task

Signal-chaining enables an application to link and load the shared library `libjsig.so` before the system libraries. The `libjsig.so` library ensures that calls such as `signal()`, `sigset()`, and `sigaction()` are intercepted so that their handlers do not replace the JVM’s signal handlers. Instead, these calls save the new signal handlers, or “chain” them behind the handlers that are installed by the JVM. Later, when any of these signals are raised and found not to be targeted at the JVM, the preinstalled handlers are invoked.

If you install signal handlers that use `sigaction()`, some `sa_flags` are not observed when the JVM uses the signal. These are:

- `SA_NOCLDSTOP` - This is always unset.
- `SA_NOCLDWAIT` - This is always unset.
- `SA_RESTART` - This is always set.

The `libjsig.so` library also hides JVM signal handlers from the application. Therefore, calls such as `signal()`, `sigset()`, and `sigaction()` that are made after the JVM has started no longer return a reference to the JVM’s signal handler, but instead return any handler that was installed before JVM startup.

The environment variable **JAVA_HOME** should be set to the location of the SDK, for example, *install_dir*.

To use libjsg.a:

- Link it with the application that creates or embeds a JVM:

```
cc_r -q64 <other compile/link parameter> -L install_dir  
-ljsig -L install_dir/jre/bin/j9vm -ljvm java_application.c
```

Note: Use `xlc_r` or `xlC_r` in place of `cc_r` if that is how you usually call the compiler or linker.

Writing JNI applications

Valid Java Native Interface (JNI) version numbers that programs can specify on the `JNI_CreateJavaVM()` API call are: `JNI_VERSION_1_2(0x00010002)` and `JNI_VERSION_1_4(0x00010004)`.

Restriction: Version 1.1 of the JNI is not supported.

This version number determines only the level of the JNI to use. The actual level of the JVM that is created is specified by the JSE libraries (use the `java -version` command to show the JVM level). The JNI level *does not* affect the language specification that is implemented by the JVM, the class library APIs, or any other area of JVM behavior. For more information, see <http://download.oracle.com/javase/6/docs/technotes/guides/jni/>.

If your application needs two JNI libraries, one built for 31- and the other for 64-bit, use the **com.ibm.vm.bitmode** system property to determine if you are running with a 31- or 64-bit JVM and choose the appropriate library.

For more information about writing 64-bit applications, see the IBM Redpaper *z/OS 64-bit C/C++ and Java Programming Environment* at <http://www.redbooks.ibm.com/abstracts/redp9110.html>.

ASCII and EBCDIC issues

On z/OS, the Java Virtual Machine is essentially an EBCDIC application. Enhanced ASCII methods are C or C++ code that has been compiled with ASCII compiler options. If you create JNI routines as enhanced ASCII C or C++ methods you will be operating in a bimodal environment; your application will be crossing over between ASCII and EBCDIC environments.

The inherent problem with bimodal programs is that, in the z/OS runtime environment, threads are designated as either EBCDIC or enhanced ASCII and are not intended to be switched between these modes in typical use. Enhanced ASCII is not designed to handle bimodal issues. You might get unexpected results or experience failures when the active mode does not match that of the compiled code. There are z/OS runtime calls that applications might use to switch the active mode between EBCDIC and enhanced ASCII (the `__ae_thread_swapmode()` and `__ae_thread_setmode()` functions are documented in Language Environment® Vendor Interfaces, see the *SA22-7568-06 Red Book*: <http://publibz.boulder.ibm.com/epubs/pdf/ceev1160.pdf>). However, even if an application is carefully coded to switch modes correctly, other bimodal issues might exist.

Supported compilers

These compilers have been tested with the IBM SDK.

The c89 compiler packaged with z/OS v1.9, C/OS/390 C++ Optional Feature is supported for:

- 31-bit z/OS on S/390®
- 64-bit z/OS on S/390

Native formatting of Java types long, double, float

The latest C/C++ compilers and runtime environments can convert jlong, jdouble, and jfloat data types to strings by using printf()-type functions.

Previous versions of the SDK for z/OS 31-bit had a set of native conversion functions and macros for formatting large Java data types. These functions and macros were:

ll2str() function

Converts a jlong to an ASCII string representation of the 64-bit value.

flt2dbl() function

Converts a jfloat to a jdouble.

dbl2nat() macro

Converts a jdouble to an ESA/390 native double.

dbl_sqrt() macro

Calculates the square root of a jdouble and returns it as a jdouble.

dbl2str() function

Converts a jdouble to an ASCII string representation.

flt2str() function

Converts a jfloat to an ASCII string representation.

These functions and macros are no longer supported by Version 6 of the SDK for z/OS. To provide a migration path, the functions have been moved to the demos area of the SDK. The demo code for these functions has been updated to reflect the changes.

The functions ll2str(), dbl2str(), and flt2str() are provided in the following object files:

- *install_dir*/demo/jni/JNINativeTypes/c/convert.o (For 31-bit)
- *install_dir*/demo/jni/JNINativeTypes/c/convert64.o (For 64-bit)

The function flt2dbl() and the macros dbl2nat() and dbl_sqrt() are not defined. However, the following macros give their definitions:

```
#include <math.h>
#define flt2dbl(f) ((double)f)
#define dbl2nat(a) ((a))
#define dbl_sqrt(a) (sqrt(a))
```

A C/C++ application that returns a jfloat data type to a Java application must be compiled with the FLOAT (IEEE) C/C++ compiler option. Applications compiled without this option return incorrect data types. Further information about compiling C/C++ source code, which applies to this Java release, can be found in the article <http://www.ibm.com/developerworks/java/library/j-jni/>

Support for thread-level recovery of blocked connectors

Four new IBM-specific SDK classes have been added to the `com.ibm.jvm` package to support the thread-level recovery of Blocked connectors. The new classes are packaged in `core.jar`.

These classes allow you to unblock threads that have become blocked on networking or synchronization calls. If an application does not use these classes, it must end the whole process, rather than interrupting an individual blocked thread.

The classes are:

public interface InterruptibleContext

Defines two methods, `isBlocked()` and `unlock()`. The other three classes implement `InterruptibleContext`.

public class InterruptibleLockContext

A utility class for interrupting synchronization calls.

public class InterruptibleIOContext

A utility class for interrupting network calls.

public class InterruptibleThread

A utility class that extends `java.lang.Thread`, to allow wrapping of interruptible methods. It uses instances of `InterruptibleLockContext` and `InterruptibleIOContext` to perform the required `isBlocked()` and `unlock()` methods depending on whether a synchronization or networking operation is blocking the thread.

Both `InterruptibleLockContext` and `InterruptibleIOContext` work by referencing the current thread. Therefore if you do not use `InterruptibleThread`, you must provide your own class that extends `java.lang.Thread`, to use these new classes.

API documentation to support the package containing these classes is available here: [API documentation](#)

Configuring large page memory allocation

You can enable large page support, on systems that support it, by starting Java with the `-Xlp` option.

About this task

Large page usage is primarily intended to provide performance improvements to applications that allocate a great deal of memory and frequently access that memory. The large page performance improvements are a result of the reduced number of misses in the Translation Lookaside Buffer (TLB). The TLB maps a larger virtual storage area range and thus causes this improvement.

For more information about the `-Xlp` options, see “JVM command-line options” on page 82.

The following restrictions apply to large page sizes on z/OS:

1M nonpageable

- This page size is supported only on the 64-bit SDK for z/OS, not the 31-bit SDK.

- This page size requires z/OS V1.9 or later with APAR OA25485, and a System z10[®] processor or later.
- A system programmer must configure z/OS for 1M nonpageable large pages.
- Users who require large pages must be authorized to the IARRSM.LRGPAGES resource in the RACF[®] (or an equivalent security product) FACILITY class with read authority.

To obtain the large page sizes available and the current setting, use the **-verbose:sizes** option. Note the current settings are the requested sizes and not the sizes obtained. For object heap size information, check the **-verbose:gc** output.

CORBA support

The Java Platform, Standard Edition (JSE) supports, at a minimum, the specifications that are defined in the compliance document from Oracle. In some cases, the IBM JSE ORB supports more recent versions of the specifications.

The minimum specifications supported are defined in the Official Specifications for CORBA support in Java SE 6: <http://download.oracle.com/javase/6/docs/api/org/omg/CORBA/doc-files/compliance.html>.

Support for GIOP 1.2

This SDK supports all versions of GIOP, as defined by chapters 13 and 15 of the CORBA 2.3.1 specification, OMG document *formal/99-10-07*.

<http://www.omg.org/cgi-bin/doc?formal/99-10-07>

Bidirectional GIOP is not supported.

Support for Portable Interceptors

This SDK supports Portable Interceptors, as defined by the OMG in the document *ptc/01-03-04*, which you can obtain from:

<http://www.omg.org/cgi-bin/doc?ptc/01-03-04>

Portable Interceptors are hooks into the ORB that ORB services can use to intercept the normal flow of execution of the ORB.

Support for Interoperable Naming Service

This SDK supports the Interoperable Naming Service, as defined by the OMG in the document *ptc/00-08-07*, which you can obtain from:

<http://www.omg.org/cgi-bin/doc?ptc/00-08-07>

The default port that is used by the Transient Name Server (the **tnameserv** command), when no **ORBInitialPort** parameter is given, has changed from *900* to *2809*, which is the port number that is registered with the IANA (Internet Assigned Number Authority) for a CORBA Naming Service. Programs that depend on this default might have to be updated to work with this version.

The initial context that is returned from the Transient Name Server is now an `org.omg.CosNaming.NamingContextExt`. Existing programs that narrow the reference to a context `org.omg.CosNaming.NamingContext` still work, and do not need to be recompiled.

The ORB supports the `-ORBInitRef` and `-ORBDefaultInitRef` parameters that are defined by the Interoperable Naming Service specification, and the `ORB::string_to_object` operation now supports the ObjectURL string formats (corbaloc: and corbaname:) that are defined by the Interoperable Naming Service specification.

The OMG specifies a method `ORB::register_initial_reference` to register a service with the Interoperable Naming Service. However, this method is not available in the Oracle Java Core API at this release. Programs that have to register a service in the current version must invoke this method on the IBM internal ORB implementation class. For example, to register a service "MyService":

```
((com.ibm.CORBA.iiop.ORB)orb).register_initial_reference("MyService",
serviceRef);
```

Where `orb` is an instance of `org.omg.CORBA.ORB`, which is returned from `ORB.init()`, and `serviceRef` is a CORBA Object, which is connected to the ORB. This mechanism is an interim one, and is not compatible with future versions or portable to non-IBM ORBs.

System properties for tracing the ORB

A runtime debug feature provides improved serviceability. You might find it useful for problem diagnosis or it might be requested by IBM service personnel.

Tracing Properties

`com.ibm.CORBA.Debug=true|fine|finer|finest`
Turns on ORB tracing.

`com.ibm.CORBA.CommTrace=true`
Adds GIOP messages (sent and received) to the trace.

`com.ibm.CORBA.Debug.Output=<file>`
Specify the trace output file. By default, this is of the form `orbtrc.DDMMYYYY.HHmm.SS.txt`.

Example of ORB tracing

For example, to trace events and formatted GIOP messages from the command line, type:

```
java -Dcom.ibm.CORBA.Debug=true
-Dcom.ibm.CORBA.CommTrace=true <myapp>
```

Limitations

Do not enable tracing for normal operation, because it might cause performance degradation. Even if you have switched off tracing, FFDC (First Failure Data Capture) is still working, so serious errors are reported. If a debug output file is generated, examine it to check on the problem. For example, the server might have stopped without performing an `ORB.shutdown()`.

The content and format of the trace output might vary from version to version.

System properties for tuning the ORB

The ORB can be tuned to work well with your specific network. The properties required to tune the ORB are described here.

com.ibm.CORBA.FragmentSize=<*size in bytes*>

Used to control GIOP 1.2 fragmentation. The default size is 1024 bytes.

To disable fragmentation, set the fragment size to 0 bytes:

```
java -Dcom.ibm.CORBA.FragmentSize=0 <myapp>
```

com.ibm.CORBA.RequestTimeout=<*time in seconds*>

Sets the maximum time to wait for a CORBA Request. By default the ORB waits indefinitely. Do not set the timeout too low to avoid connections ending unnecessarily.

com.ibm.CORBA.LocateRequestTimeout=<*time in seconds*>

Set the maximum time to wait for a CORBA LocateRequest. By default the ORB waits indefinitely.

com.ibm.CORBA.ListenerPort=<*port number*>

Set the port for the ORB to read incoming requests on. If this property is set, the ORB starts listening as soon as it is initialized. Otherwise, it starts listening only when required.

Java security permissions for the ORB

When running with a Java SecurityManager, invocation of some methods in the CORBA API classes might cause permission checks to be made, which might result in a SecurityException. If your program uses any of these methods, ensure that it is granted the necessary permissions.

Table 6. Methods affected when running with Java SecurityManager

Class/Interface	Method	Required permission
org.omg.CORBA.ORB	init	java.net.SocketPermission resolve
org.omg.CORBA.ORB	connect	java.net.SocketPermission listen
org.omg.CORBA.ORB	resolve_initial_references	java.net.SocketPermission connect
org.omg.CORBA. portable.ObjectImpl	_is_a	java.net.SocketPermission connect
org.omg.CORBA. portable.ObjectImpl	_non_existent	java.net.SocketPermission connect
org.omg.CORBA. portable.ObjectImpl	OutputStream _request (String, boolean)	java.net.SocketPermission connect
org.omg.CORBA. portable.ObjectImpl	_get_interface_def	java.net.SocketPermission connect
org.omg.CORBA. Request	invoke	java.net.SocketPermission connect
org.omg.CORBA. Request	send_deferred	java.net.SocketPermission connect
org.omg.CORBA. Request	send_oneway	java.net.SocketPermission connect
javax.rmi. PortableRemoteObject	narrow	java.net.SocketPermission connect

ORB implementation classes

A list of the ORB implementation classes.

The ORB implementation classes in this release are:

- `org.omg.CORBA.ORBClass=com.ibm.CORBA.iiop.ORB`
- `org.omg.CORBA.ORBSingletonClass=com.ibm.rmi.corba.ORBSingleton`
- `javax.rmi.CORBA.UtilClass=com.ibm.CORBA.iiop.UtilDelegateImpl`
- `javax.rmi.CORBA.StubClass=com.ibm.rmi.javax.rmi.CORBA.StubDelegateImpl`
- `javax.rmi.CORBA.PortableRemoteObjectClass`
`=com.ibm.rmi.javax.rmi.PortableRemoteObject`

These are the default values, and you are advised not to set these properties or refer to the implementation classes directly. For portability, make references only to the CORBA API classes, and not to the implementation. These values might be changed in future releases.

RMI over IIOP

Java Remote Method Invocation (RMI) provides a simple mechanism for distributed Java programming. RMI over IIOP (RMI-IIOP) uses the Common Object Request Broker Architecture (CORBA) standard Internet Inter-ORB Protocol (IIOP) to extend the base Java RMI to perform communication. This allows direct interaction with any other CORBA Object Request Brokers (ORBs), whether they were implemented in Java or another programming language.

The following documentation is available:

- The RMI-IIOP Programmer's Guide is an introduction to writing RMI-IIOP programs.
- The *Java Language to IDL Mapping* document is a detailed technical specification of RMI-IIOP: <http://www.omg.org/cgi-bin/doc?ptc/00-01-06.pdf>.

Implementing the Connection Handler Pool for RMI

Thread pooling for RMI Connection Handlers is not enabled by default.

About this task

To enable the connection pooling implemented at the RMI TCPTransport level, set the option

```
-Dsun.rmi.transport.tcp.connectionPool=true
```

This version of the Runtime Environment does not have a setting that you can use to limit the number of threads in the connection pool.

Enhanced BigDecimal

From Java 5.0, the IBM `BigDecimal` class has been adopted by Oracle as `java.math.BigDecimal`. The `com.ibm.math.BigDecimal` class is reserved for possible future use by IBM and is currently deprecated. Migrate existing Java code to use `java.math.BigDecimal`.

The new `java.math.BigDecimal` uses the same methods as both the previous `java.math.BigDecimal` and `com.ibm.math.BigDecimal`. Existing code using `java.math.BigDecimal` continues to work correctly. The two classes do not serialize.

To migrate existing Java code to use the `java.math.BigDecimal` class, change the import statement at the start of your `.java` file from: `import com.ibm.math.*;` to `import java.math.*;`

Working in a multiple network stack environment

In a multiple network stack environment (CINET), when one of the stacks fails, no notification or Java exception occurs for a Java program that is listening on an `INADDR_ANY` socket. Also, when new stacks become available, the Java application does not become aware of them until it rebinds the `INADDR` socket.

To avoid this situation, when a TCP/IP stack comes online:

- If the `ibm.socketserver.recover` property is set to `false` (which is the default), an exception (`NetworkRecycledException`) is thrown to the application to allow it either to fail or to attempt to rebind.
- If the `ibm.socketserver.recover` property is set to `true`, Java attempts to redrive the socket connection on the new stack if listening on all addresses (`addrs`). If the socket bind cannot be replayed at that time, an exception (`NetworkRecycledException`) is thrown to the application to allow it either to fail or to attempt to rebind.

Both `ServerSocket.accept()` and `ServerSocketChannel.accept()` can throw `NetworkRecycledException`.

While a socket is listening for new connections, it maintains a queue of incoming connections. When `NetworkRecycledException` is thrown and the system attempts to rebind the socket, the connection queue is reset and connection requests in this queue are dropped.

Support for XToolkit

XToolkit is included by default. You need XToolkit when using the `SWT_AWT` bridge in Eclipse to build an application that uses both SWT and Swing.

Restriction: Motif is no longer supported and will be removed in a later release.

Related links:

- An example of integrating Swing into Eclipse RCPs: <http://eclipsezone.com/eclipse/forums/t45697.html>
- Reference Information in the Eclipse information center: http://help.eclipse.org/help32/index.jsp?topic=/org.eclipse.platform.doc.isv/reference/api/org/eclipse/swt/awt/SWT_AWT.html
- Set up information is available on the Oracle Corporation Web site: <http://download.oracle.com/javase/6/docs/technotes/guides/awt/1.5/xawt.html>

Support for the Java Attach API

Your application can connect to another “target” virtual machine using the Java Attach API. Your application can then load an agent application into the target virtual machine, for example to perform tasks such as monitoring status. Support for the Java Attach API was added in Java 6 SR 6.

Code for agent applications, such as JMX agents or JVM TI agents, is normally loaded during virtual machine startup by specifying special startup parameters. Requiring startup parameters might not be convenient for using agents on applications that are already running, such as WebSphere Application Servers. You can use the Java Attach API to load an agent at any time, by specifying the process ID of the target virtual machine. The Attach API capability is sometimes called the “late attach” capability.

Support for the Attach API is disabled by default for Java 6 SR 6 and later. For Java 6 SR 9 and later, processes that use the default z/OS OMVS segment cannot enable the attach API for security reasons.

Security considerations

Security for the Java Attach API is handled by POSIX file permissions. On z/OS, you must use UNIX user permissions to protect your applications. It is not sufficient to rely on RACF or system level security to protect your applications. The reason is that these mechanisms do not have the necessary UNIX permissions set up and configured for the Java Attach API to remain secure.

The Java Attach API creates files and directories in a common directory.

The key security features of the Java Attach API are:

- A process using the Java Attach API must be owned by the same UNIX user ID as the target process. This constraint ensures that only the target process owner or root can attach other applications to the target process.
- The common directory uses the sticky bit to prevent a user from deleting or replacing a subdirectory belonging to another user. To preserve the security of this mechanism, set the ownership of the common directory to ROOT. This directory will contain files such as `_attachlock`, `_master`, and `_notifier`, which are used only for synchronization. These files can be owned by any user, and must have read and write permission. However, you can remove execute permission on these files, if present. The files are empty and will be recreated automatically if deleted.
- The files in the subdirectory for a process, with the exception of a lock file, are accessible only by the owner of a process. The subdirectory has owner read, write, and execute permissions plus group and world execute permissions. In this directory, read and write access are restricted to the owner only, except for the `attachNotificationSync` file, which must have world and group write permissions. This exception does not affect security because the file is used exclusively for synchronization and is never written to or read.
- Information about the target process can be written and read only by the owner.
- Java 5 SR10 allowed users in the same group to access to each others' processes. This capability was removed in later versions.

You must secure access to the Java Attach API capability to ensure that only authorized users or processes can connect to another virtual machine. If you do not intend to use the Java Attach API capability, disable this feature using a Java system property. Set the `com.ibm.tools.attach.enable` system property to the value `no`; for example:

```
-Dcom.ibm.tools.attach.enable=no
```

The Attach API can be enabled by setting the `com.ibm.tools.attach.enable` system property to the value `yes`; for example:

```
-Dcom.ibm.tools.attach.enable=yes
```

Using the Java Attach API

By default, the target virtual machine is identified by its process ID. To use a different target, change the system property **com.ibm.tools.attach.id**; for example:

```
-Dcom.ibm.tools.attach.id=<process_ID>
```

The target process also has a human-readable “display name”. By default, the display name is the command line used to start Java. To change the default display name, use the **com.ibm.tools.attach.displayName** system property. The ID and display name cannot be changed after the application has started.

The Attach API creates working files in a common directory, which by default is called `.com_ibm_tools_attach` and is created in the system temporary directory. The system property **java.io.tmpdir** holds the value of the system temporary directory. On non-Windows systems, the system temporary directory is typically `/tmp`.

You can specify a different common directory from the default, by using the following Java system property:

```
-Dcom.ibm.tools.attach.directory=directory_name
```

This system property causes the specified directory, *directory_name*, to be used as the common directory. If the directory does not already exist, it is created, however the parent directory must already exist. For example, the following system property creates a common directory called `myattachapidir` in the `usr` directory. The `usr` directory must already exist.

```
-Dcom.ibm.tools.attach.directory=/usr/myattachapidir
```

The common directory must be located on a local drive; specifying a network mounted file system might result in incorrect behavior.

If your Java application ends abnormally, for example, following a crash or a SIGKILL signal, the process subdirectory is not deleted. The Java VM detects and removes obsolete subdirectories where possible. The subdirectory can also be deleted by the owning user ID.

On heavily loaded system, applications might experience timeouts when attempting to connect to target applications. The default timeout is 120 seconds. Use the **com.ibm.tools.attach.timeout** system property to specify a different timeout value in milliseconds. For example, to timeout after 60 seconds:

```
-Dcom.ibm.tools.attach.timeout=60000
```

A timeout value of zero indicates an indefinite wait.

For JMX applications, you can disable authentication by editing the `<JAVA_HOME>/jre/lib/management/management.properties` file. Set the following properties to disable authentication in JMX:

```
com.sun.management.jmxremote.authenticate=false  
com.sun.management.jmxremote.ssl=false
```

Problems with the Attach API result in one of the following exceptions:

- `com.ibm.tools.attach.AgentLoadException`
- `com.ibm.tools.attach.AgentInitializationException`

- `com.ibm.tools.attach.AgentNotSupportedException`
- `java.io.IOException`

A useful reference for information about the Attach API can be found at <http://download.oracle.com/javase/6/docs/technotes/guides/attach/index.html>. The IBM implementation of the Attach API is equivalent to the Oracle Corporation implementation. However, the IBM implementation cannot be used to attach to, or accept attach requests from, non-IBM virtual machines. To use the attach API to attach to target processes from your application, you must add the "tools.jar" library to the application classpath. This library is not required for the target processes to accept attach requests.

Chapter 6. Applet Viewer

The `appletviewer` is used to test applications designed to be run in a browser.

Distributing Java applications

Java applications typically consist of class, resource, and data files.

When you distribute a Java application, your software package probably consists of the following parts:

- Your own class, resource, and data files
- An installation procedure or program

Your SDK for z/OS software license does **not** allow you to redistribute any of the SDK's files with your application. You must ensure that a licensed version of the SDK for z/OS is installed on the target workstation.

When distributing your application for use on a z/OS platform, make the z/OS SDK a prerequisite, because z/OS does not have a separate JRE.

Chapter 7. Class data sharing between JVMs

Class data sharing enables multiple JVMs to share a single space in memory.

You can share class data between Java Virtual Machines (JVMs) by storing it in a cache in shared memory. Sharing reduces the overall virtual storage consumption when more than one JVM shares a cache. Sharing also reduces the startup time for a JVM after the cache has been created. The shared class cache is independent of any running JVM and persists until it is deleted or until a system IPL.

A shared cache can contain:

- Bootstrap classes
- Application classes
- Metadata that describes the classes
- Ahead-of-time (AOT) compiled code

| The format of classes stored in the shared classes cache is changed in IBM SDK for
| z/OS, V6 service refresh 13. As a result, there is a different shared cache generation
| number, which causes the JVM to create a new shared classes cache, rather than
| re-creating or reusing an existing cache. To save space, all existing shared caches
| should be removed unless they are in use by an earlier release of IBM SDK for
| z/OS, V6. For more information about deleting a shared classes cache, see “Class
| data sharing command-line options” on page 57.

Overview of class data sharing

Class data sharing provides a method of reducing memory footprint and improving JVM start time. Java 6 provides new and improved features in cache management, isolation, and performance.

Enabling class data sharing

Enable class data sharing by using the **-Xshareclasses** option when starting a JVM. The JVM connects to an existing cache or creates a new cache if one does not exist.

All bootstrap and application classes loaded by the JVM are shared by default. Custom class loaders share classes automatically if they extend the application class loader. Otherwise, they must use the Java Helper API provided with the JVM to access the cache. See “Adapting custom class loaders to share classes” on page 65.

The JVM can also store ahead-of-time (AOT) compiled code in the cache for certain methods to improve the startup time of subsequent JVMs. The AOT compiled code is not shared between JVMs, but is cached to reduce compilation time when the JVM starts. The amount of AOT code stored in the cache is determined heuristically. You cannot control which methods get stored in the cache. You can set maximum and minimum limits on the amount of cache space used for AOT code, or you can disable AOT caching completely. See “Class data sharing command-line options” on page 57 for more information.

Cache access

A JVM can access a cache with either read/write or read-only access. Any JVM connected to a cache with read/write access can update the cache. Any number of JVMs can concurrently read from the cache, even while another JVM is writing to it.

You must take care if runtime bytecode modification is being used. See “Runtime bytecode modification” on page 63 for more information.

Dynamic updating of the cache

The shared class cache persists beyond the lifetime of any JVM. Therefore, the cache is updated dynamically to reflect any modifications that might have been made to JARs or classes on the file system. The dynamic updating makes the cache independent of the application using it.

Cache security

Access to the shared class cache is limited by operating system permissions and Java security permissions. The shared class cache is created with user access by default unless the **groupAccess** command-line suboption is used. Only a class loader that has registered to share class data can update the shared class cache.

A cache can be accessed only by a JVM running in the same storage key as the JVM that created the cache. If the keys do not match, permission to access the cache is denied. Known environments where storage keys can cause an issue include:

WebSphere control region (key 2)

Attempting to access the shared cache for the WebSphere control region generates the following error message:

```
JVMSHRC337W Platform error message: shmat : EDC5111I Permission denied.
```

CICS, when switching between **STGPROT=NO** (key 8) and **STGPROT=YES** (key 9)

If CICS is started with **STGPROT=YES**, CICS allocates the shared cache in key 9. This enables cache sharing between programs running in CICS key (8) and User key (9). If CICS is started with **STGPROT=NO**, the cache is allocated in key 8. Using a shared cache in key 8 might lead to errors if the CICS region is subsequently restarted with **STGPROT=YES**. Unless the cache is rebuilt, a program running in User key (9) will be unable to access the shared class cache. In this situation the JVM issues a message, similar to:

```
JVMSHRC337W Platform error message: shmat : EDC5111I Permission denied.
```

The default storage key for the JVM is key 8.

(31-bit only) The cache memory is protected against accidental or deliberate corruption using memory page protection. This protection is not an absolute guarantee against corruption because the JVM must unprotect pages to write to them. The only way to guarantee that a cache cannot be modified is to open it read-only.

(64-bit only) Memory page protection using PGSER PROTECT is unavailable on z/OS 64-bit mode.

If a Java SecurityManager is installed, classloaders, excluding the default bootstrap, application, and extension class loaders, must be granted permission to share

classes. Grant permission by adding `SharedClassPermission` lines to the `java.policy` file. See “Using `SharedClassPermission`” on page 64 for more information. The `RuntimePermission createClassLoader` restricts the creation of new class loaders and therefore also restricts access to the cache.

Cache lifespan

Multiple caches can exist on a system and you specify them by name as a suboption to the `-Xshareclasses` command. A JVM can connect to only one cache at any one time.

You can override the default cache size on startup using `-Xscmx<n><size>`. This size is then fixed for the lifetime of the cache. Caches exist until they are explicitly deleted using a suboption to the `-Xshareclasses` command or until the next system IPL.

Cache utilities

All cache utilities are suboptions to the `-Xshareclasses` command. See “Class data sharing command-line options” or use `-Xshareclasses:help` to see a list of available suboptions.

Class data sharing command-line options

Class data sharing and the cache management utilities are controlled using command-line options to the Java launcher.

For options that take a `<size>` parameter, suffix the number with "k" or "K" to indicate kilobytes, "m" or "M" to indicate megabytes, or "g" or "G" to indicate gigabytes.

`-Xscmaxaot<size>`

Sets the maximum number of bytes in the cache that can be used for AOT data. Use this option to ensure that a certain amount of cache space is available for non-AOT data. By default, the maximum limit for AOT data is the amount of free space in the cache. The value of this option should not be smaller than the value of `-Xscminaot` and must not be larger than the value of `-Xscmx`.

`-Xscminaot<size>`

Sets the minimum number of bytes in the cache to reserve for AOT data. By default, no space is reserved for AOT data, although AOT data is written to the cache until the cache is full or the `-Xscmaxaot` limit is reached. The value of this option must not exceed the value of `-Xscmx` or `-Xscmaxaot`. The value of `-Xscminaot` must always be considerably less than the total cache size because AOT data can be created only for cached classes. If the value of `-Xscminaot` is equal to the value of `-Xscmx`, no class data or AOT data is stored because AOT data must be associated with a class in the cache.

`-Xscmx<size>`

Specifies cache size. This option applies only if a cache is being created and no cache of the same name exists. The default cache size is platform-dependent. You can find out the size value being used by adding `-verbose:sizes` as a command-line argument. The minimum cache size is 4 KB. The maximum cache size is also platform-dependent. (See “Cache size limits” on page 62.)

`-Xshareclasses:<suboption>[,<suboption>...]`

Enables class data sharing. Can take a number of suboptions, some of which

are cache utilities. Cache utilities perform the required operation on the specified cache, without starting the VM. You can combine multiple suboptions, separated by commas, but the cache utilities are mutually exclusive. When running cache utilities, the message Could not create the Java virtual machine is expected. Cache utilities do not create the virtual machine.

Some cache utilities can work with caches from previous Java versions or caches that are created by JVMs with different bit-widths. These caches are referred to as “incompatible” caches.

You can use the following suboptions with the **-Xshareclasses** option:

help

Lists all the command-line suboptions.

name=<name>

Connects to a cache of a given name, creating the cache if it does not already exist. Also used to indicate the cache that is to be modified by cache utilities; for example, **destroy**. Use the **listAllCaches** utility to show which named caches are currently available. If you do not specify a name, the default name “sharedcc_%u” is used. %u in the cache name inserts the current user name. You can specify “%g” in the cache name to insert the current group name.

Note: Some features, when enabled, result in the creation of caches that cannot be shared with caches that were created when the feature was disabled. The multitenancy support is one such example. In this situation, you can have more than one cache with the same name.

cacheDir=<directory>

Sets the directory in which cache data is read and written. By default, <directory> is /tmp/javasharedresources. The user must have sufficient permissions in <directory>. Caches are stored in shared memory and have control files that describe the location of the memory. Control files are stored in a javasharedresources subdirectory of the **cacheDir** specified. Do not move or delete control files in this directory. The **listAllCaches** utility, the **destroyAll** utility, and the **expire** suboption work only in the scope of a given **cacheDir**.

cacheDirPerm=<permission>

Sets UNIX-style permissions when creating a cache directory. <permission> must be a number in the ranges 0700 - 0777 or 1700 - 1777. If <permission> is not valid, the JVM terminates with an appropriate error message.

The permissions that are specified by this suboption are used only when creating a new cache directory. If the cache directory already exists, this suboption is ignored and the cache directory permissions are not changed.

If you set this suboption to 0000, the default directory permissions are used. If you set this suboption to 1000, the machine default directory permissions are used, but the sticky bit is enabled.

If the cache directory is the platform default directory, /tmp/javasharedresources, the **cacheDirPerm** suboption is ignored and the cache directory permissions are set to 777. If you do not set the **cacheDirPerm** suboption, and the cache directory does not already exist, a new directory is created with permissions set to 777, for compatibility with

earlier Java versions. Permissions for existing cache directories are unchanged, to avoid generating RACF errors, which generate log messages.

readonly

Opens an existing cache with read-only permissions. The JVM does not create a new cache with this suboption. Opening a cache read-only prevents the JVM from making any updates to the cache. It also allows the JVM to connect to caches created by other users or groups without requiring write access. By default, this suboption is not specified.

groupAccess

Sets operating system permissions on a new cache to allow group access to the cache. Group access can be set only when permitted by the operating system **umask** setting. The default is user access only.

verbose

Enables verbose output, which provides overall status on the shared class cache and more detailed error messages.

verboseAOT

Enables verbose output when compiled AOT code is being found or stored in the cache. AOT code is generated heuristically. You might not see any AOT code generated at all for a small application. You can disable AOT caching by using the **noaot** suboption.

verboseIO

Gives detailed output on the cache I/O activity, listing information on classes that are stored and found. Each class loader is given a unique ID (the bootstrap loader is always 0) and the output shows the class loader hierarchy at work, where class loaders must ask their parents for a class before they can load it themselves. It is usual to see many failed requests; this behavior is expected for the class loader hierarchy.

verboseHelper

Enables verbose output for the Java Helper API. This output shows you how the Helper API is used by your class loader.

silent

Turns off all shared classes messages, including error messages. Unrecoverable error messages, which prevent the JVM from initializing, are displayed.

nonfatal

Allows the JVM to start even if class data sharing fails. Normal behavior for the JVM is to refuse to start if class data sharing fails. If you select **nonfatal** and the shared classes cache fails to initialize, the JVM attempts to connect to the cache in read-only mode. If this attempt fails, the JVM starts without class data sharing.

none

Can be added to the end of a command line to disable class data sharing. This suboption overrides class sharing arguments found earlier on the command line.

modified=<modified context>

Used when a JVMTI agent is installed that might modify bytecode at run time. If you do not specify this suboption and a bytecode modification agent is installed, classes are safely shared with an extra performance cost. The *<modified context>* is a descriptor that is chosen by the user; for example, "myModification1". This option partitions the cache, so that only

JVMs that use context myModification1 can share the same classes. For instance, if you run HelloWorld with a modification context and then run it again with a different modification context, all classes are stored twice in the cache. For more information, see “Runtime bytecode modification” on page 63.

reset

Causes a cache to be destroyed and then recreated when the JVM starts up. Can be added to the end of a command line as **-Xshareclasses:reset**.

destroy (Utility option)

Destroys a cache specified by the **name**, **cacheDir**, and **nonpersistent** suboptions. A cache can be destroyed only if all JVMs using it have shut down, and the user has sufficient permissions.

destroyAll (Utility option)

Tries to destroy all caches available using the specified **cacheDir** and **nonpersistent** suboptions. A cache can be destroyed only if all JVMs using it have shut down, and the user has sufficient permissions.

expire=<time in minutes>

Destroys all caches that have been unused for the time that is specified before loading shared classes. This option is not a utility option because it does not cause the JVM to exit.

listAllCaches (Utility option)

Lists all the compatible and incompatible caches that exist in the specified cache directory. If you do not specify **cacheDir**, the default directory is used. Summary information, such as Java version and current usage is displayed for each cache.

Note: Some features, when enabled, result in the creation of caches that cannot be shared with caches that are created when the feature is disabled. The multitenancy support is one such example. In this situation, you can have more than one cache with the same name. The output from the **listAllCaches** option has a feature column which lists the feature that created the cache, usually default. For multitenancy support, the feature is mt, and the cache is listed in the Incompatible shared caches section of the output.

printStats (Utility option)

Displays summary information for the cache that is specified by the **name**, **cacheDir**, and **nonpersistent** suboptions. The most useful information that is displayed is how full the cache is and how many classes it contains. Stale classes are classes that are updated on the file system and which the cache has therefore marked as "stale". Stale classes are not purged from the cache and can be reused. For more information, see printStats utility.

printAllStats (Utility option)

Displays detailed information for the cache that is specified by the **name**, **cacheDir**, and **nonpersistent** suboptions. Every class is listed in chronological order, with a reference to the location from which it was loaded. AOT code for class methods is also listed.

For more information, see printAllStats utility.

(31-bit only) mprotect=[all | default | none]

By default, the memory pages that contain the cache are always protected, unless a specific page is being updated. This protection helps prevent accidental or deliberate corruption to the cache. The cache header is not

protected by default because this protection has a small performance cost. Specifying `all` ensures that all the cache pages are protected, including the header. Specifying `none` disables the page protection.

noBootclasspath

Prevents storage of classes that are loaded by the bootstrap class loader in the shared classes cache. Can be used with the `SharedClassURLFilter` API to control exactly which classes get cached. For more information about shared class filtering, see [Using the SharedClassHelper API](#).

cacheRetransformed

Enables caching of classes that are transformed by using the `JVMTI RetransformClasses` function.

noaot

Disables caching of AOT code. AOT code already in the shared data cache can be loaded.

Creating, populating, monitoring, and deleting a cache

An overview of the life-cycle of a shared class data cache including examples of the cache management utilities.

To enable class data sharing, add `-Xshareclasses[:name=<name>]` to your application command line.

The JVM either connects to an existing cache of the given name or creates a new cache of that name. If a new cache is created, it is populated with all bootstrap and application classes being loaded until the cache becomes full. If two or more JVMs are started concurrently, they populate the cache concurrently.

To check that the cache has been created, run `java -Xshareclasses:listAllCaches`. To see how many classes and how much class data is being shared, run `java -Xshareclasses:[name=<name>],printStats`. You can run these utilities after the application JVM has terminated or in another command window.

Note: Some features, when enabled, result in the creation of caches that cannot be shared with caches that are created when the feature is disabled. The multitenancy support is one such example. In this situation, you can have more than one cache with the same name. The output from the `listAllCaches` option has a feature column which lists the feature that created the cache, usually `default`. For multitenancy support, the feature is `mt`, and the cache is listed in the `Incompatible shared caches` section of the output.

For more feedback on cache usage while the JVM is running, use the **verbose** suboption. For example, `java -Xshareclasses:[name=<name>],verbose`.

To see classes being loaded from the cache or stored in the cache, add `-Xshareclasses:[name=<name>],verboseIO` to your application command line.

To delete the cache, run `java -Xshareclasses:[name=<name>],destroy`. You usually delete caches only if they contain many stale classes or if the cache is full and you want to create a bigger cache.

You should tune the cache size for your specific application, because the default is unlikely to be the optimum size. To determine the optimum cache size, specify a large cache, using `-Xscmx`, run the application, and then use **printStats** to

determine how much class data has been stored. Add a small amount to the value shown in **printStats** for contingency. Because classes can be loaded at any time during the lifetime of the JVM, it is best to do this analysis after the application has terminated. However, a full cache does not have a negative affect on the performance or capability of any JVMs connected to it, so it is acceptable to decide on a cache size that is smaller than required.

If a cache becomes full, a message is displayed on the command line of any JVMs using the **verbose** suboption. All JVMs sharing the full cache then loads any further classes into their own process memory. Classes in a full cache can still be shared, but a full cache is read-only and cannot be updated with new classes.

Performance and memory consumption

Class data sharing is particularly useful on systems that use more than one JVM running similar code; the system benefits from reduced real storage consumption. It is also useful on systems that frequently start and shut down JVMs, which benefit from the improvement in startup time.

The processor and memory usage required to create and populate a new cache is minimal. The JVM startup cost in time for a single JVM is typically between 0 and 5% slower compared with a system not using class data sharing, depending on how many classes are loaded. JVM startup time improvement with a populated cache is typically between 10% and 40% faster compared with a system not using class data sharing, depending on the operating system and the number of classes loaded. Multiple JVMs running concurrently show greater overall startup time benefits.

Duplicate classes are consolidated in the shared class cache. For example, class A loaded from `myClasses.jar` and class A loaded from `myOtherClasses.jar` (with identical content) is stored only once in the cache. The **printAllStats** utility shows multiple entries for duplicated classes, with each entry pointing to the same class.

When you run your application with class data sharing, you can use the operating system tools to see the reduction in virtual storage consumption.

Considerations and limitations of using class data sharing

Consider these factors when deploying class data sharing in a product and using class data sharing in a development environment.

Cache size limits

The maximum theoretical cache size is 2 GB. The size of cache you can specify is limited by the amount of physical memory and swap space available to the system.

Because the virtual address space of a process is shared between the shared classes cache and the Java heap, if you increase the maximum size of the Java heap you might reduce the size of the shared classes cache you can create.

JVMTI `RetransformClasses()` is unsupported

You cannot run `RetransformClasses()` on classes loaded from the shared class cache.

The JVM might throw the exception `UnmodifiableClassException` if you attempt to run `RetransformClasses()`. It does not work because class file bytes are not available for classes loaded from the shared class cache. If you must use `RetransformClasses()`, ensure that the classes to be transformed are not loaded from the shared class cache, or disable the shared class cache feature.

Required APAR for Shared Classes

You must apply z/OS APAR OA11519, available for z/OS R1.6 and onwards, to any z/OS system where shared classes are used. This APAR ensures that multiple `shmat` requests for the same shared segment will map to the same virtual address for multiple JVMs.

Without this APAR, there is a problem with using shared memory when multiple JVMs are stored in a single address space. Each `shmat` call consumes a separate virtual address range. This is not acceptable because shared classes will run out of shared memory pages prematurely.

Working with BPXPRMxx settings

Some of the **BPXPRMxx** parmlib settings affect shared classes performance. Using the wrong settings can stop shared classes from working. These settings might also have performance implications.

For further information about performance implications, and use of these parameters, see the *z/OS MVS Initialization and Tuning Reference (SA22-7592)* at <http://publibz.boulder.ibm.com/epubs/pdf/iea2e280.pdf> and the *z/OS Unix System Services Planning Guide (GA22-7800)* at <http://publibz.boulder.ibm.com/epubs/pdf/bpxzb280.pdf>. The most significant **BPXPRMxx** parameters that affect the operation of shared classes are:

- **MAXSHAREPAGES**, **IPCSHMSPAGES**, **IPCSHMMPAGES**, and **IPCSHMNSEGS**. These settings affect the amount of shared memory pages available to the JVM. The JVM uses these memory pages for the shared classes cache. If you request large cache sizes, you might have to increase the amount of shared memory pages available. The shared page size for a z/OS UNIX System Service is fixed at 4 KB for 31-bit and 1 MB for 64-bit. Shared classes try to create a 16 MB cache by default on both 31- and 64-bit platforms. Therefore set **IPCSHMMPAGES** greater than 4096 on a 31-bit system.
If you set a cache size with **-Xscmx**, the VM rounds up the value to the nearest megabyte. You must take this factoring into account when setting **IPCSHMMPAGES** on your system.
- **IPCSEMNIIDS**, and **IPCSEMNSEMS**. These settings affect the amount of SystemV IPC semaphore available to UNIX processes. IBM shared classes use System V IPC semaphores to communicate between the JVMs.

Runtime bytecode modification

Any JVM using a JVM Tool Interface (JVMTI) agent that can modify bytecode data must use the **modified=<modified_context>** suboption if it wants to share the modified classes with another JVM.

The modified context is a user-specified descriptor that describes the type of modification being performed. The modified context partitions the cache so that all JVMs running under the same context share a partition.

This partitioning allows JVMs that are not using modified bytecode to safely share a cache with those that are using modified bytecode. All JVMs using a given modified context must modify bytecode in a predictable, repeatable manner for each class, so that the modified classes stored in the cache have the expected modifications when they are loaded by another JVM. Any modification must be predictable because classes loaded from the shared class cache cannot be modified again by the agent.

If a JVMTI agent is used without a modification context, classes are still safely shared by the JVM, but with a small affect on performance. Using a modification context with a JVMTI agent avoids the need for extra checks and therefore has no affect on performance. A custom `ClassLoader` that extends `java.net.URLClassLoader` and modifies bytecode at load time without using JVMTI automatically stores that modified bytecode in the cache, but the cache does not treat the bytecode as modified. Any other VM sharing that cache loads the modified classes. You can use the `modified=<modification_context>` suboption in the same way as with JVMTI agents to partition modified bytecode in the cache. If a custom `ClassLoader` needs to make unpredictable load-time modifications to classes, that `ClassLoader` must not attempt to use class data sharing.

See [Dealing with runtime bytecode modification](#) for more detail on this topic.

Operating system limitations

Temporary disk space must be available to hold cache information. The operating system enforces cache permissions.

The shared class cache requires disk space to store identification information about the caches that exist on the system. This information is stored in `/tmp/javasharedresources`. If the identification information directory is deleted, the JVM cannot identify the shared classes on the system and must re-create the cache. Use the `ipcs` command to view the memory segments used by a JVM or application.

Users running a JVM must be in the same group to use a shared class cache. The operating system enforces the permissions for accessing a shared class cache. If you do not specify a cache name, the user name is appended to the default name so that multiple users on the same system create their own caches by default.

Using SharedClassPermission

If a `SecurityManager` is being used with class data sharing and the running application uses its own class loaders, you must grant these class loaders shared class permissions before they can share classes.

You add shared class permissions to the `java.policy` file using the `ClassLoader` class name (wildcards are permitted) and either `"read"`, `"write"`, or `"read,write"` to determine the access granted. For example:

```
permission com.ibm.oti.shared.SharedClassPermission
    "com.abc.customclassloaders.*", "read,write";
```

If a `ClassLoader` does not have the correct permissions, it is prevented from sharing classes. You cannot change the permissions of the default bootstrap, application, or extension class loaders.

Adapting custom class loaders to share classes

Any class loader that extends `java.net.URLClassLoader` can share classes without modification. You must adopt class loaders that do not extend `java.net.URLClassLoader` to share class data.

You must grant all custom class loaders shared class permissions if a `SecurityManager` is being used; see “Using `SharedClassPermission`” on page 64. IBM provides several Java interfaces for various types of custom class loaders, which allow the class loaders to find and store classes in the shared class cache. These classes are in the `com.ibm.oti.shared` package.

The API documentation for this package is available here: [API documentation](#)

See [Using the Java Helper API](#) for more information about how to use these interfaces.

Chapter 8. Service and support for independent software vendors

Contact points for service:

If you are entitled to services for the Program code pursuant to the IBM Solutions Developer Program, contact the IBM Solutions Developer Program through your usual method of access or on the Web at: <http://www.ibm.com/partnerworld/>.

If you have purchased a service contract (that is, the IBM Personal Systems Support Line or equivalent service by country), the terms and conditions of that service contract determine what services, if any, you are entitled to receive with respect to the Program.

Chapter 9. Accessibility

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

IBM strives to provide products with usable access for everyone, regardless of age or ability.

For example, you can operate the IBM SDK for z/OS, V6 without a mouse, by using only the keyboard.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

For users who require keyboard navigation, a description of useful keystrokes for Swing applications can be found here: [Swing Key Bindings](#).

IBM and accessibility

See the IBM Human Ability and Accessibility Center for more information about the commitment that IBM has to accessibility.

Keyboard traversal of JComboBox components in Swing

If you traverse the drop-down list of a JComboBox component with the cursor keys, the button or editable field of the JComboBox does not change value until an item is selected. This is the correct behavior for this release and improves accessibility and usability by ensuring that the keyboard traversal behavior is consistent with mouse traversal behavior.

Chapter 10. Any comments on this user guide?

If you have any comments about this user guide, contact us through one of the following channels. Note that these channels are not set up to answer technical queries, but are for comments about the documentation only.

Send your comments:

- By e-mail to idrcf@hursley.ibm.com.
- By fax:
 - From the UK: 01962 842327
 - From elsewhere: +44 1962 842327
- By mail to:

IBM United Kingdom Ltd
User Technologies,
Mail Point 095
Hursley Park
Winchester
Hampshire
SO21 2JN
United Kingdom

The fine print. By choosing to send a message to IBM, you acknowledge that all information contained in your message, including feedback data, such as questions, comments, suggestions, or the like, shall be deemed to be non-confidential and IBM shall have no obligation of any kind with respect to such information and shall be free to reproduce, use, disclose, and distribute the information to others without limitation. Further, IBM shall be free to use any ideas, concepts, know-how or techniques contained in such information for any purpose whatsoever, including, but not limited to, developing, manufacturing and marketing products incorporating such information.

Appendix. Appendixes

Reference information.

Command-line options

You can specify the options on the command line while you are starting Java. They override any relevant environment variables. For example, using `-cp <dir1>` with the Java command completely overrides setting the environment variable `CLASSPATH=<dir2>`.

This chapter provides the following information:

- “Specifying command-line options”
- “General command-line options” on page 74
- “System property command-line options” on page 75
- “JVM command-line options” on page 82
- “JIT and AOT command-line options” on page 98
- “Garbage Collector command-line options” on page 101

Specifying command-line options

Although the command line is the traditional way to specify command-line options, you can also pass options to the JVM by using options files and environment variables.

The sequence of the Java options on the command line defines which options take precedence during startup. Rightmost options have precedence over leftmost options. In the following example, the `-Xjit` option takes precedence:

```
java -Xint -Xjit myClass
```

Use single or double quotation marks for command-line options only when explicitly directed to do so. Single and double quotation marks have different meanings on different platforms, operating systems, and shells. Do not use `'-X<option>'` or `"-X<option>"`. Instead, you must use `-X<option>`. For example, do not use `'-Xmx500m'` and `"-Xmx500m"`. Write this option as `-Xmx500m`.

At startup, the list of JVM arguments is constructed in the following order, with the lowest precedence first:

1. Environment variables that are described in `../com.ibm.java.doc.diagnostics.60/diag/appendixes/env_var/env_jvm.dita` are translated into command-line options. For example, the following environment variable adds the parameter `-Xrs` to the list of arguments:

```
export IBM_NOSIGHANDLER=<non_null_string>
```
2. The `IBM_JAVA_OPTIONS` environment variable. You can set command-line options using this environment variable. The options that you specify with this environment variable are added to the command line when a JVM starts in that environment.

The environment variable can contain multiple blank-delimited argument strings, but must not contain comments. For example:

```
export IBM_JAVA_OPTIONS="-Dmysysprop1=tcip -Dmysysprop2=wait -Xdisablejavadump"
```

Note: The environment variable **JAVA_TOOLS_OPTIONS** is equivalent to **IBM_JAVA_OPTIONS** and is available for compatibility with JVM TI.

3. Certain options are created automatically by the JVM. These specify arguments such as search paths and version information.

4. Options that are specified on the command line. For example:

```
java -Dmysysprop1=tcPIP -Dmysysprop2=wait -Xdisablejvadm MyJavaClass
```

The Java launcher adds some automatically generated arguments to this list, such as the names of the main class.

You can also use the **-Xoptionsfile** parameter to specify JVM options. This parameter can be used on the command line, or as part of the **IBM_JAVA_OPTIONS** environment variable. The contents of an option file are expanded in place during startup. For more information about the structure and contents of this type of file, see “-Xoptionsfile” on page 89.

To troubleshoot startup problems, you can check which options are used by a JVM. Append the following command-line option, and inspect the Jvadm file that is generated:

```
-Xdump:java:events=vmstart
```

Here is an extract from a Jvadm file that shows the options that are used:

```
....
2CIUSERARG          -Xdump:java:file=/home/test_javacore.txt,events=vmstop
2CIUSERARG          -Dtest.cmdlineOption=1
2CIUSERARG          -XXallowvmshutdown:true
2CIUSERARG          -Xoptionsfile=test1.test_options_file
....
```

General command-line options

Use these options to print help on assert-related options, set the search path for application classes and resources, print a usage method, identify memory leaks inside the JVM, print the product version and continue, enable verbose output, and print the product version.

**-cp, -classpath <directories and compressed or .jar files separated by : (
on Windows)>**

Sets the search path for application classes and resources. If **-classpath** and **-cp** are not used, and the **CLASSPATH** environment variable is not set, the user classpath is, by default, the current directory (.).

-help, -?

Prints a usage message.

-fullversion

Prints the build and version information for the JVM.

-showversion

Prints product version and continues.

-verbose:<option>[,<option>...]

Enables verbose output. Separate multiple options using commas. These options are available:

class

Writes an entry to stderr for each class that is loaded.

dynload

Provides detailed information as each bootstrap class is loaded by the JVM:

- The class name and package

- For class files that were in a .jar file, the name and directory path of the .jar
- Details of the size of the class and the time taken to load the class

The data is written out to stderr. An example of the output on a Windows platform follows:

```
<Loaded java/lang/String from C:\sdk\jre\lib\vm.jar>
<Class size 17258; ROM size 21080; debug size 0>
<Read time 27368 usec; Load time 782 usec; Translate time 927 usec>
```

gc Provide verbose garbage collection information.

init

Writes information to stderr describing JVM initialization and termination.

jni

Writes information to stderr describing the JNI services called by the application and JVM.

sizes

Writes information to stderr describing the active memory usage settings.

stack

Writes information to stderr describing the Java and C stack usage for each thread.

-version

Prints the full build and version information for the JVM.

System property command-line options

Use the system property command-line options to set up your system.

-D<name>=<value>

Sets a system property.

-Dcom.ibm.dbgmalloc

This option provides memory allocation diagnostic information for class library native code.

-Dcom.ibm.dbgmalloc=true

When an application is started with this option, a javadump records the amount of memory allocated by the class library components. You can use this option together with the **-Xcheck:memory** option to obtain information about class library call sites and their allocation sizes. Enabling this option has an impact on throughput performance. The information does not include allocation information from Abstract Windows Toolkit (AWT), ZLIB data compression library, and libwrapper ASCII to EBCDIC conversion library.

-Dcom.ibm.jsse2.renegotiate

If your Java application uses JSSE for secure communication, you can disable TLS renegotiation by installing APAR IZ65239.

-Dcom.ibm.jsse2.renegotiate=[ALL | NONE | ABBREVIATED]

ALL Allow both abbreviated and unabbreviated (full) renegotiation handshakes.

NONE

Allow no renegotiation handshakes. This value is the default setting.

ABBREVIATED

Allow only abbreviated renegotiation handshakes.

-Dcom.ibm.lang.management.verbose

Enables verbose information from java.lang.management operations to be written to the output channel during VM operation.

-Dcom.ibm.lang.management.verbose

There are no options for this system property.

-Dcom.ibm.IgnoreMalformedInput

From Java 6 SR9, any invalid UTF8 or malformed byte sequences are replaced with the standard unicode replacement character `\uFFFD`.

-Dcom.ibm.IgnoreMalformedInput=true

To retain the old behavior, where invalid UTF8 or malformed byte sequences are ignored, set this system property to *true*.

-Dcom.ibm.streams.CloseFDWithStream

Determines whether the `close()` method of a stream object closes a native file descriptor even if the descriptor is still in use by another stream object.

-Dcom.ibm.streams.CloseFDWithStream=[true | false]

Usually, you create a `FileInputStream` or `FileOutputStream` instance by passing a `String` or a `File` object to the stream constructor method. Each stream then has a separate file descriptor. However, you can also create a stream by using an existing `FileDescriptor` instance, for example one that you obtain from a `RandomAccessFile` instance, or another `FileInputStream` or `FileOutputStream` instance. Multiple streams can then share the same file descriptor.

If you set this option to *false*, when you use the `close()` method of the stream, the associated file descriptor is also closed only if it is not in use by any other streams. If you set the option to *true*, the file descriptor is closed regardless of any other streams that might still be using it.

The default setting is *true*.

Note: Before version 6 service refresh 14, the default behavior was to close the file descriptor only when all the streams that were using it were also closed. This system property exists so that you can revert to this previous default behavior if necessary. This system property will be removed in a future release, so you should adjust your applications to use the new default behavior before you upgrade to a later release.

-Dcom.ibm.tools.attach.enable

Enable the Attach API for this application.

-Dcom.ibm.tools.attach.enable=yes

The Attach API allows your application to connect to a virtual machine. Your application can then load an agent application into the virtual machine. The agent can be used to perform tasks such as monitoring the virtual machine status.

-Dcom.ibm.UseCLDR16

This property reverts behavior to an earlier releases.

-Dcom.ibm.UseCLDR16

From IBM SDK for z/OS, V6 V6 service refresh 10, changes are made to the locale translation files to make them consistent with Oracle JDK 6. To understand the differences in detail, see <http://www.ibm.com/support/docview.wss?uid=swg21568667>. Include the **-Dcom.ibm.UseCLDR16** system property on the command-line to revert to the locale translation files used in earlier releases.

-Dcom.ibm.xtq.processor.overrideSecureProcessing

This system property affects the XSLT processing of extension functions or extension elements when Java security is enabled.

Purpose

From IBM SDK for z/OS, V6 service refresh 14, the use of extension functions or extension elements is not allowed when Java security is enabled. This change is introduced to enhance security. The system property can be used to revert to the behavior in earlier releases.

Parameters

com.ibm.xtq.processor.overrideSecureProcessing=true

To revert to the behavior in earlier releases of the IBM SDK, set this system property to *true*.

-Dcom.ibm.zipfile.closeinputstreams

The Java.util.zip.ZipFile class allows you to create InputStreams on files held in a compressed archive.

-Dcom.ibm.zipfile.closeinputstreams=true

Under some conditions, using ZipFile.close() to close all InputStreams that have been opened on the compressed archive might result in a 56-byte-per-InputStream native memory leak. Setting the **-Dcom.ibm.zipfile.closeinputstreams=true** forces the JVM to track and close InputStreams without the memory impact caused by retaining native-backed objects. Native-backed objects are objects that are stored in native memory, rather than the Java heap. By default, the value of this system property is not enabled.

-Dfile.encoding

Use this property to define the file encoding that is required.

-Dfile.encoding=value

Where *value* defines the file encoding that is required.

By default the IBM GBK converter follows Unicode 3.0 standards. To force the IBM GBK converter to follow Unicode 2.0 standards, use a value of *bestfit936*.

-Dibm.jvm.bootclasspath

The value of this property is used as an additional search path.

-Dibm.jvm.bootclasspath

The value of this property is used as an additional search path, which is inserted between any value that is defined by **-Xbootclasspath/p:** and the bootclass path. The bootclass path is either the default or the one that you defined by using the **-Xbootclasspath:** option.

-Dibm.stream.nio

This option addresses the ordering of IO and NIO converters.

-Dibm.stream.nio=[true | false]

When this option is set to *true*, the NIO converters are used instead of the IO converters. By default the IO converters are used.

-Djava.compiler

Disables the Java compiler by setting to NONE.

-Djava.compiler=[NONE | j9jit<vm_version>]

Enable JIT compilation by setting to `j9jit<vm_version>` (Equivalent to `-Xjit`).

-Djavax.xml.namespace.QName.useCompatibleHashCodeAlgorithm

Use this property to turn off an enhanced hashing algorithm for `javax.xml.namespace.QName.hashCode()`.

-Djavax.xml.namespace.QName.useCompatibleHashCodeAlgorithm=1.0

From Java 6 SR11 an enhanced hashing algorithm is used for `javax.xml.namespace.QName.hashCode()`. This algorithm can change the iteration order of items returned from hash maps. For compatibility, you can restore the earlier hashing algorithm by setting the system property

-Djavax.xml.namespace.QName.useCompatibleHashCodeAlgorithm=1.0.

-Djdk.map.althashing.threshold

This system property controls the use of an enhanced hashing algorithm for hashed maps.

-Djdk.map.althashing.threshold=*value*

This alternative hashing algorithm is used for string keys when a hashed data structure has a capacity larger than *value*.

A value of *1* ensures that this algorithm is always used, regardless of the hashed map capacity. A value of *-1* prevents the use of this algorithm, which is the default value.

The hashed map structures affected by this threshold are: `java.util.HashMap`, `java.util.Hashtable`, `java.util.LinkedHashMap`, `java.util.WeakHashMap`, and `java.util.concurrent.ConcurrentHashMap`.

The capacity of a hashed map is related to the number of entries in the map, multiplied by the load factor. Because the capacity of a hashed map is rounded up to the next power of two, setting the threshold to intermediate values has no affect on behavior. For example, threshold values of 600, 700, and 1000 have the same effect. However, values of 1023 and 1024 cause a difference in behavior. For a more detailed description of the capacity and load factor, see <http://docs.oracle.com/javase/7/docs/api/java/util/HashMap.html>.

When entries are removed from a hashed map the capacity does not shrink. Therefore, if the map ever exceeds the threshold to use alternative hashing for Strings, the map always uses alternative hashing for Strings. This behavior does not change, even if entries are later removed or the map is emptied using `clear()`.

The enhanced hashing algorithm is available from Java 6 SR11

-Djdk.xml.entityExpansionLimit

This option provides limits for Java API for XML (JAXP) processing. Use this option to limit the number of entity expansions in an XML document.

-Djdk.xml.entityExpansionLimit=*value*

Where *value* is a positive integer. The default value is 64,000.

A value of *0* or a negative number sets no limits.

You can also set this limit by adding the following line to your `jaxp.properties` file:

`jdk.xml.entityExpansionLimit=<value>`

| **Related reference:**

| “-Djdk.xml.maxGeneralEntitySizeLimit”

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the maximum size of a general entity.

| “-Djdk.xml.maxOccur”

| This option provides limits for Java API for XML (JAXP) processing. This option
| defines the maximum number of content model nodes that can be created in a
| grammar.

| “-Djdk.xml.maxParameterEntitySizeLimit” on page 80

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the maximum size of a parameter entity.

| “-Djdk.xml.totalEntitySizeLimit” on page 81

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the total size of all entities that include general and parameter
| entities.

| **-Djdk.xml.maxGeneralEntitySizeLimit**

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the maximum size of a general entity.

| To protect an application from malformed XML, set this value to the minimum size
| possible.

| **-Djdk.xml.maxGeneralEntitySizeLimit=*value***

| Where *value* is the maximum size that is allowed for a general entity. The
| default value is 0.

| A value of 0 or a negative number sets no limits.

| You can also set this limit by adding the following line to your `jaxp.properties`
| file:

| `jdk.xml.maxGeneralEntitySizeLimit=<value>`

| **Related reference:**

| “-Djdk.xml.entityExpansionLimit” on page 78

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the number of entity expansions in an XML document.

| “-Djdk.xml.maxOccur”

| This option provides limits for Java API for XML (JAXP) processing. This option
| defines the maximum number of content model nodes that can be created in a
| grammar.

| “-Djdk.xml.maxParameterEntitySizeLimit” on page 80

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the maximum size of a parameter entity.

| “-Djdk.xml.totalEntitySizeLimit” on page 81

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the total size of all entities that include general and parameter
| entities.

| **-Djdk.xml.maxOccur**

| This option provides limits for Java API for XML (JAXP) processing. This option
| defines the maximum number of content model nodes that can be created in a
| grammar.

When building a grammar for a W3C XML schema, use this option to limit the number of content model nodes that can be created when the schema defines attributes that can occur multiple times.

-Djdk.xml.maxOccur=*value*

Where *value* is a positive integer. The default value is 5,000.

A value of 0 or a negative number sets no limits.

You can also set this limit by adding the following line to your `jaxp.properties` file:

```
jdk.xml.maxoccur=<value>
```

Related reference:

“-Djdk.xml.entityExpansionLimit” on page 78

This option provides limits for Java API for XML (JAXP) processing. Use this option to limit the number of entity expansions in an XML document.

“-Djdk.xml.maxGeneralEntitySizeLimit” on page 79

This option provides limits for Java API for XML (JAXP) processing. Use this option to limit the maximum size of a general entity.

“-Djdk.xml.maxParameterEntitySizeLimit”

This option provides limits for Java API for XML (JAXP) processing. Use this option to limit the maximum size of a parameter entity.

“-Djdk.xml.totalEntitySizeLimit” on page 81

This option provides limits for Java API for XML (JAXP) processing. Use this option to limit the total size of all entities that include general and parameter entities.

-Djdk.xml.maxParameterEntitySizeLimit

This option provides limits for Java API for XML (JAXP) processing. Use this option to limit the maximum size of a parameter entity.

To protect an application from malformed XML, set this value to the minimum size possible.

-Djdk.xml.maxParameterEntitySizeLimit=*value*

Where *value* is the maximum size that is allowed for a parameter entity. The default value is 0.

A value of 0 or a negative number sets no limits.

You can also set this limit by adding the following line to your `jaxp.properties` file:

```
jdk.xml.maxParameterEntitySizeLimit=<value>
```


| **Related reference:**

| “-Djdk.xml.entityExpansionLimit” on page 78

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the number of entity expansions in an XML document.

| “-Djdk.xml.maxGeneralEntitySizeLimit” on page 79

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the maximum size of a general entity.

| “-Djdk.xml.maxOccur” on page 79

| This option provides limits for Java API for XML (JAXP) processing. This option
| defines the maximum number of content model nodes that can be created in a
| grammar.

| “-Djdk.xml.totalEntitySizeLimit”

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the total size of all entities that include general and parameter
| entities.

| **-Djdk.xml.totalEntitySizeLimit**

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the total size of all entities that include general and parameter
| entities.

| **-Djdk.xml.totalEntitySizeLimit=*value***

| Where *value* is the collective size of all entities. The default value is 5×10^7 .

| A value of 0 or a negative number sets no limits.

| You can also set this limit by adding the following line to your `jaxp.properties`
| file:

| `jdk.xml.totalEntitySizeLimit=<value>`

| **Related reference:**

| “-Djdk.xml.entityExpansionLimit” on page 78

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the number of entity expansions in an XML document.

| “-Djdk.xml.maxGeneralEntitySizeLimit” on page 79

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the maximum size of a general entity.

| “-Djdk.xml.maxOccur” on page 79

| This option provides limits for Java API for XML (JAXP) processing. This option
| defines the maximum number of content model nodes that can be created in a
| grammar.

| “-Djdk.xml.maxParameterEntitySizeLimit” on page 80

| This option provides limits for Java API for XML (JAXP) processing. Use this
| option to limit the maximum size of a parameter entity.

| **-Dsun.awt.keepWorkingSetOnMinimize**

| The **-Dsun.awt.keepWorkingSetOnMinimize=true** system property stops the JVM
| trimming an application when it is minimized.

| **-Dsun.awt.keepWorkingSetOnMinimize=true**

| When a Java application using the Abstract Windowing Toolkit (AWT) is
| minimized, the default behavior is to “trim” the “working set”. The working
| set is the application memory stored in RAM. Trimming means that the
| working set is marked as being available for swapping out if the memory is
| required by another application. The advantage of trimming is that memory is

available for other applications. The disadvantage is that a “trimmed” application might experience a delay as the working set memory is brought back into RAM.

The default behavior is to trim an application when it is minimized.

-Dsun.net.client.defaultConnectTimeout

Specifies the default value for the connect timeout for the protocol handlers used by the `java.net.URLConnection` class.

-Dsun.net.client.defaultConnectTimeout=<value in milliseconds>

The default value set by the protocol handlers is -1, which means that no timeout is set.

When a connection is made by an applet to a server and the server does not respond properly, the applet might seem to hang. The delay might also cause the browser to hang. The apparent hang occurs because there is no network connection timeout. To avoid this problem, the Java Plug-in has added a default value to the network timeout of 2 minutes for all HTTP connections. You can override the default by setting this property.

-Dsun.net.client.defaultReadTimeout

Specifies the default value for the read timeout for the protocol handlers used by the `java.net.URLConnection` class when reading from an input stream when a connection is established to a resource.

-Dsun.net.client.defaultReadTimeout=<value in milliseconds>

The default value set by the protocol handlers is -1, which means that no timeout is set.

-Dsun.nio.MaxDirectMemorySize

Limits the native memory size for `nio` Direct Byte Buffer objects to the value specified.

-Dsun.nio.MaxDirectMemorySize=<value>

Specify *<value>* in bytes.

-Dsun.rmi.transport.tcp.connectionPool

Enables thread pooling for the RMI ConnectionHandlers in the TCP transport layer implementation.

-Dsun.rmi.transport.tcp.connectionPool=*val*

val is either *true* or a value that is not null.

-Dswing.useSystemFontSettings

This option addresses compatibility problems for Swing programs.

-Dswing.useSystemFontSettings=[*false*]

By default, Swing programs running with the Windows Look and Feel render with the system font set by the user instead of a Java-defined font. As a result, fonts differ from the fonts in earlier releases. This option addresses compatibility problems like these for programs that depend on the old behavior. By setting this option, v1.4.1 fonts and those of earlier releases are the same for Swing programs running with the Windows Look and Feel.

JVM command-line options

Use these options to configure your JVM. The options prefixed with **-X** are nonstandard.

Options that relate to the JIT are listed under “JIT and AOT command-line options” on page 98. Options that relate to the Garbage Collector are listed under “Garbage Collector command-line options” on page 101.

-X

Displays help on nonstandard options.

-X Displays help on nonstandard options.

-Xaggressive

Enables performance optimizations.

-Xaggressive

Enables performance optimizations that are expected to be the default in future releases.

-Xargencoding

Include Unicode escape sequences in the argument list.

-Xargencoding

You can put Unicode escape sequences in the argument list. This option is set to off by default.

-Xbootclasspath

Sets the search path for bootstrap classes and resources.

-Xbootclasspath:*<directories and compressed or Java archive files separated by : (; on Windows)>*

The default is to search for bootstrap classes and resources in the internal VM directories and .jar files.

-Xbootclasspath/a:

Appends to the end of the search path for bootstrap classes.

-Xbootclasspath/a:*<directories and compressed or Java archive files separated by : (; on Windows)>*

Appends the specified directories, compressed files, or .jar files to the end of the bootstrap class path. The default is to search for bootstrap classes and resources in the internal VM directories and .jar files.

-Xbootclasspath/p:

Adds a prefix to the search path for bootstrap classes.

-Xbootclasspath/p:*<directories and compressed or Java archive files separated by : (; on Windows)>*

Adds a prefix of the specified directories, compressed files, or Java archive files to the front of the bootstrap class path. Do not deploy applications that use the **-Xbootclasspath:** or the **-Xbootclasspath/p:** option to override a class in the standard API. The reason is that such a deployment contravenes the Java 2 Runtime Environment binary code license. The default is to search for bootstrap classes and resources in the internal VM directories and .jar files.

-Xcheck

You can use the **-Xcheck** option to run checks during JVM startup, such as memory checks or checks on JNI functions.

-Xcheck:*<option>*

The options available are detailed in separate topics.

-Xcheck:classpath:

Displays a warning message if an error is discovered in the class path.

-Xcheck:classpath

Checks the classpath and reports if an error is discovered; for example, a missing directory or JAR file.

-Xcheck:gc:

Runs additional checks on garbage collection.

-Xcheck:gc[:<scan options>][:<verify options>][:<misc options>]

By default, no checks are made. See the output of **-Xcheck:gc:help** for more information.

-Xcheck:jni:

Runs additional checks for JNI functions.

-Xcheck:jni[:help][:<option>=<value>]

This option is equivalent to **-Xrunjnichk**. By default, no checks are made.

-Xcheck:memory:

Identifies memory leaks inside the JVM.

-Xcheck:memory[:<option>]

Identifies memory leaks inside the JVM using strict checks that cause the JVM to exit on failure. If no option is specified, **all** is used by default. The available options are as follows:

all

Enables checking of all allocated and freed blocks on every free and allocate call. This check of the heap is the most thorough. It typically causes the JVM to exit on nearly all memory-related problems soon after they are caused. This option has the greatest affect on performance.

callsite=<number of allocations>

Displays callsite information every *<number of allocations>*. De-allocations are not counted. Callsite information is presented in a table with separate information for each callsite. Statistics include:

- The number and size of allocation and free requests since the last report.
- The number of the allocation request responsible for the largest allocation from each site.

Callsites are presented as *sourcefile:linenumber* for C code and assembly function name for assembler code.

Callsites that do not provide callsite information are accumulated into an "unknown" entry.

failat=<number of allocations>

Causes memory allocation to fail (return NULL) after *<number of allocations>*. Setting *<number of allocations>* to 13 causes the 14th allocation to return NULL. De-allocations are not counted. Use this option to ensure that JVM code reliably handles allocation failures. This option is useful for checking allocation site behavior rather than setting a specific allocation limit.

ignoreUnknownBlocks

Ignores attempts to free memory that was not allocated using the **-Xcheck:memory** tool. Instead, the **-Xcheck:memory** statistics that are printed out at the end of a run indicates the number of “unknown” blocks that were freed.

mprotect=<top|bottom>

Locks pages of memory on supported platforms, causing the program to stop if padding before or after the allocated block is accessed for reads or writes. An extra page is locked on each side of the block returned to the user.

If you do not request an exact multiple of one page of memory, a region on one side of your memory is not locked. The **top** and **bottom** options control which side of the memory area is locked. **top** aligns your memory blocks to the top of the page (lower address), so buffer underruns result in an application failure. **bottom** aligns your memory blocks to the bottom of the page (higher address) so buffer overruns result in an application failure.

Standard padding scans detect buffer underruns when using **top** and buffer overruns when using **bottom**.

nofree

Keeps a list of blocks that are already used instead of freeing memory. This list, and the list of currently allocated blocks, is checked for memory corruption on every allocation and deallocation. Use this option to detect a dangling pointer (a pointer that is “dereferenced” after its target memory is freed). This option cannot be reliably used with long-running applications (such as WebSphere Application Server), because “freed” memory is never reused or released by the JVM.

noscan

Checks for blocks that are not freed. This option has little effect on performance, but memory corruption is not detected. This option is compatible only with **subAllocator**, **callsite**, and **callsitesmall**.

quick

Enables block padding only and is used to detect basic heap corruption. Every allocated block is padded with sentinel bytes, which are verified on every allocate and free. Block padding is faster than the default of checking every block, but is not as effective.

skipto=<number of allocations>

Causes the program to check only on allocations that occur after *<number of allocations>*. De-allocations are not counted. Use this option to speed up JVM startup when early allocations are not causing the memory problem. The JVM performs approximately 250+ allocations during startup.

subAllocator[=<size in MB>]

Allocates a dedicated and contiguous region of memory for all JVM allocations. This option helps to determine if user JNI code or the JVM is responsible for memory corruption. Corruption in the JVM **subAllocator** heap suggests that the JVM is causing the problem; corruption in the user-allocated memory suggests that user code is corrupting memory. Typically, user and JVM allocated memory are interleaved.

zero

Newly allocated blocks are set to 0 instead of being filled with the 0xE7E7xxxxxxxE7E7 pattern. Setting these blocks to 0 helps you to

determine whether a callsite is expecting zeroed memory, in which case the allocation request is followed by `memset(pointer, 0, size)`.

Note: The `-Xcheck:memory` option cannot be used in the `-Xoptionsfile`.

-Xclassgc

Enables dynamic unloading of classes by the JVM. Garbage collection of class objects occurs only on class loader changes.

-Xclassgc

Dynamic unloading is the default behavior. To disable dynamic class unloading, use the `-Xnoclassgc` option.

-Xcompressedrefs

Enables the use of compressed references.

-Xcompressedrefs

(64-bit only) To disable compressed references, use the `-Xnocompressedreferences` option. For more information, see Compressed references.

Compressed references are disabled by default.

You cannot include this option in an options file. You must specify this option on the command line, or by using the `IBM_JAVA_OPTIONS` environment variable.

-Xdbg

Loads debugging libraries to support the remote debugging of applications.

-Xdbg:<options>

This option is deprecated in the IBM SDK for z/OS, V6. By default, the debugging libraries are not loaded, and the VM instance is not enabled for debug.

The preferred method to enable the debugger is `-agentlib:jdwp=<options>`. For more information about using the Java debugger, see the IBM SDK for z/OS, V6 user guide.

-Xdiagnosticscollector

Enables the Diagnostics Collector.

-Xdiagnosticscollector[:settings=<filename>]

See The Diagnostics Collector for more information. The settings option allows you to specify a different Diagnostics Collector settings file to use instead of the default `dc.properties` file in the JRE.

-Xdisablejavadump

Turns off Javadump generation on errors and signals.

-Xdisablejavadump

By default, Javadump generation is enabled.

-Xdump

Use the `-Xdump` option to add and remove dump agents for various JVM events, update default dump settings (such as the dump name), and limit the number of dumps that are produced.

-Xdump

See Using dump agents for more information.

-Xenableexplicitgc

This options tells the VM to trigger a garbage collection when a call is made to `System.gc()`.

-Xenableexplicitgc

Signals to the VM that calls to `System.gc()` trigger a garbage collection. This option is enabled by default.

-Xfastresolve

Tune performance by improving the resolution time for classes.

-Xfastresolve<n>

This option is used to tune performance by improving the resolution time for classes when the field count exceeds the threshold specified by `<n>`. If profiling tools show significant costs in field resolution, change the threshold until the costs are reduced. If you enable this option, additional memory is used when the threshold is exceeded.

-Xfuture

Turns on strict class-file format checks.

-Xfuture

Use this flag when you are developing new code because stricter checks will become the default in future releases. By default, strict format checks are disabled.

-Xifa

Enables Java applications to run on IFAs if they are available.

-Xifa:<on | off | force> (z/OS only)

z/OS V1R6 or later can run Java applications on a new type of special-purpose assist processor called the *System z Application Assist Processor* (zAAP). The zAAP is also known as an IFA (Integrated Facility for Applications).

The **-Xifa** option enables Java applications to run on IFAs if they are available. The default value for the **-Xifa** option is *on*. Only Java code and system native methods can be on IFA processors.

The *force* option is obsolete and should not be used. This option is superseded by the **SYS1.PARMLIB(IEAOPTxx) PROJECTCPU=YES** parameter, which is available on all supported levels of z/OS. **Xifa:force** can be used for testing purposes when a zAAP is not available, but can have a negative performance impact.

-Xiss

Sets the initial stack size for Java threads.

-Xiss<size>

By default, the stack size is set to 2 KB. Use the **-verbose:sizes** option to output the value that the VM is using.

-Xjarversion

Produces output information about the version of each .jar file.

-Xjarversion

Produces output information about the version of each .jar file in the class path, the boot class path, and the extensions directory. Version information is taken from the Implementation-Version and Build-Level properties in the manifest of the .jar file.

Note: The **-Xjarversion** option cannot be used in the **-Xoptionsfile**.

-Xjni

Sets JNI options.

-Xjni:<suboptions>

You can use the following suboption with the **-Xjni** option:

-Xjni:arrayCacheMax=[<size in bytes>|unlimited]

Sets the maximum size of the array cache. The default size is 8096 bytes.

-Xlinenumbers

Displays line numbers in stack traces for debugging.

-Xlinenumbers

See also **-Xno1inenumbers**. By default, line numbers are on.

-Xlog

Enables message logging.

-Xlog

To prevent message logging, use the **-Xlog:none** option. By default, logging is enabled. This option is available from Java 6 SR5. See http://publib.boulder.ibm.com/infocenter/javasdk/v6r0/topic/com.ibm.java.doc.diagnostics.60/diag/appendixes/messages_generic.html.

-Xlp

Requests the JVM to allocate the Java object heap with large pages.

-Xlp[<size>]

z/OS: Requests the JVM to allocate the Java object heap using large page sizes. If <size> is not specified, the 1M nonpageable size is used. If large pages are not supported by the hardware, or enabled in RACF, the JVM does not start and produces an error message.

Allocating large pages using **-Xlp[<size>]** is only supported on the 64-bit SDK for z/OS, not the 31-bit JVM for z/OS.

On z/OS, **-Xlp<size>** only supports a large page size of 1M (nonpageable).

For more information, see “Configuring large page memory allocation” on page 44.

All platforms: To obtain the large page sizes available and the current setting, use the **-verbose:sizes** option. Note the current settings are the requested sizes and not the sizes obtained. For object heap size information, check the **-verbose:gc** output.

There is no error message issued when the operating system does not have sufficient resources to satisfy the request. This limitation and a workaround for verifying the page size used can be found in Known limitations.

-Xmso

Sets the initial stack size for operating system threads.

-Xmso<size>

The default value can be determined by running the command:

```
java -verbose:sizes
```

The maximum value for the stack size varies according to platform and specific machine configuration. If you exceed the maximum value, a `java/lang/OutOfMemoryError` message is reported.

-Xnoagent

Disables support for the old JDB debugger.

-Xnoagent

Disables support for the old JDB debugger.

-Xnoclassgc

Disables class garbage collection.

-Xnoclassgc

This option switches off garbage collection of storage associated with Java classes that are no longer being used by the JVM. The default behavior is as defined by **-Xclassgc**. Enabling this option is not recommended except under the direction of the IBM Java support team. The reason is the option can cause unlimited native memory growth, leading to out-of-memory errors.

-Xnocompressedrefs

Disables the use of compressed references.

-Xnocompressedrefs

(64-bit only) To enable compressed references, use the **-Xcompressedreferences** option. For more information, see Compressed references.

Compressed references are disabled by default.

You cannot include this option in an options file. You must specify this option on the command line, or by using the **IBM_JAVA_OPTIONS** environment variable.

-Xnolinenumbers

Disables the line numbers for debugging.

-Xnolinenumbers

See also **-Xlinenumbers**. By default, line number are on.

-Xnosigcatch

Disables JVM signal handling code.

-Xnosigcatch

See also **-Xsigcatch**. By default, signal handling is enabled.

-Xnosigchain

Disables signal handler chaining.

-Xnosigchain

See also **-Xsigchain**. By default, the signal handler chaining is enabled, except for z/OS.

-Xoptionsfile

Specifies a file that contains JVM options and definitions.

-Xoptionsfile=<file>

where <file> contains options that are processed as if they had been entered directly as command-line options. By default, a user option file is not used.

Here is an example of an options file:

```
#My options file
-X<option1>
-X<option2>=\
<value1>,\
<value2>
-D<sysprop1>=<value1>
```

The options file does not support these options:

- **-assert**
- **-fullversion**
- **-help**
- **-showversion**
- **-version**
- **-Xcompressedrefs**
- **-Xcheck:memory**
- **-Xjarversion**
- **-Xoptionsfile**

Although you cannot use **-Xoptionsfile** recursively within an options file, you can use **-Xoptionsfile** multiple times on the same command line to load more than one options files.

Some options use quoted strings as parameters. Do not split quoted strings over multiple lines using the line continuation character '\'. The '¥' character is not supported as a line continuation character. For example, the following example is not valid in an options file:

```
-Xevents=vmstop,exec="cmd /c \  
echo %pid has finished."
```

The following example is valid in an options file:

```
-Xevents=vmstop, \  
exec="cmd /c echo %pid has finished."
```

-Xoss

Sets the maximum Java stack size for any thread.

-Xoss<size>

Recognized but deprecated. Use **-Xss** and **-Xms0** instead. The maximum value for the stack size varies according to platform and specific machine configuration. If you exceed the maximum value, a java/lang/OutOfMemoryError message is reported.

-Xrdbginfo

Loads the remote debug information server with the specified host and port.

-Xrdbginfo:<host>:<port>

By default, the remote debug information server is disabled.

-Xrs

Disables signal handling in the JVM.

-Xrs

Setting **-Xrs** prevents the Java run time environment from handling any internally or externally generated signals such as SIGSEGV and SIGABRT. Any signals raised are handled by the default operating system handlers. Disabling signal handling in the JVM reduces performance by approximately 2-4%, depending on the application.

-Xrs:sync

As with **-Xrs**, the use of **-Xrs:sync** reduces performance by approximately 2-4%, depending on the application.

-Xrun

This option loads helper libraries, but has been superseded by the **-agentlib** option.

-Xrun<library name>[:<options>]

This option has been superseded; use the **-agentlib** option instead. For more information about **-agentlib**, see Using the JVM TI.

-Xrun loads helper libraries. To load multiple libraries, specify it more than once on the command line. Examples of these libraries are:

-Xrunhprof[:help] | [[:<option>=<value>, ...]

Performs heap, CPU, or monitor profiling.

-Xrunjdp[:help] | [[:<option>=<value>, ...]

Loads debugging libraries to support the remote debugging of applications. This option is the same as **-Xdbg**.

-Xrunjnichk[:help] | [[:<option>=<value>, ...]

Deprecated. Use **-Xcheck:jni** instead.

-Xscmx

Specifies cache size.

-Xscmx<size>

This option applies only if a cache is being created and no cache of the same name exists. The default cache size is platform-dependent. You can find out the size value being used by adding **-verbose:sizes** as a command-line argument. Minimum cache size is 4 KB. Maximum cache size is platform-dependent. The size of cache that you can specify is limited by the amount of physical memory and paging space available to the system. The virtual address space of a process is shared between the shared classes cache and the Java heap. Increasing the maximum size of the Java heap reduces the size of the shared classes cache that you can create.

-XselectiveDebug

Enables selective debugging.

-XselectiveDebug

Use the `com.ibm.jvm.Debuggable` annotation to mark classes and methods that must be available for debugging. The JVM optimizes methods that do not need debugging to provide better performance in a debugging environment. See the *User Guide* for your platform for more information.

-Xshareclasses

Enables class sharing.

-Xshareclasses:<suboptions>

This option can take a number of suboptions, some of which are cache utilities. Cache utilities perform the required operation on the specified cache, without starting the VM. You can combine multiple suboptions, separated by commas, but the cache utilities are mutually exclusive.

You can use the following suboptions with the **-Xshareclasses** option:

cacheDir=<directory>

Sets the directory in which cache data is read and written. By default, <directory> is `/tmp/javasharedresources` on Linux, AIX, z/OS, and IBM i. You must have sufficient permissions in <directory>. Nonpersistent caches are stored in shared memory and have control files that describe the location of the memory. Control files are stored in a `javasharedresources`

subdirectory of the **cacheDir** specified. Do not move or delete control files in this directory. The **listAllCaches** utility, the **destroyAll** utility, and the **expire** suboption work only in the scope of a given **cacheDir**.

cacheDirPerm=<permission>

Sets UNIX-style permissions when creating a cache directory. *<permission>* must be an octal number in the ranges 0700 - 0777 or 1700 - 1777. If *<permission>* is not valid, the JVM terminates with an appropriate error message.

The permissions specified by this suboption are used only when creating a new cache directory. If the cache directory already exists, this suboption is ignored and the cache directory permissions are not changed.

If you set this suboption to 0000, the default directory permissions are used. If you set this suboption to 1000, the machine default directory permissions are used, but the sticky bit is enabled. If the cache directory is the platform default directory, `/tmp/javasharedresources`, this suboption is ignored and the cache directory permissions are set to 777. If you do not set this suboption, the cache directory permissions are set to 777, for compatibility with earlier Java versions.

cacheRetransformed

Enables caching of classes that are transformed by using the JVMTI `RetransformClasses` function. See JVMTI redefinition and retransformation of classes for more information.

destroy (Utility option)

Destroys a cache that is specified by the **name**, **cacheDir**, and **nonpersistent** suboptions. A cache can be destroyed only if all JVMs using it have shut down and the user has sufficient permissions.

destroyAll (Utility option)

Tries to destroy all caches available using the specified **cacheDir** and **nonpersistent** suboptions. A cache can be destroyed only if all JVMs using it have shut down and the user has sufficient permissions.

expire=<time in minutes> (Utility option)

Destroys all caches that are unused for the time that is specified before loading shared classes. This option is not a utility option because it does not cause the JVM to exit.

groupAccess

Sets operating system permissions on a new cache to allow group access to the cache. Group access can be set only when permitted by the operating system **umask** setting. The default is user access only.

help

Lists all the command-line options.

listAllCaches (Utility option)

Lists all the compatible and incompatible caches that exist in the specified cache directory. If you do not specify **cacheDir**, the default directory is used. Summary information, such as Java version and current usage, is displayed for each cache.

Note: Some features, when enabled, result in the creation of caches that cannot be shared with caches that are created when the feature is disabled. The multitenancy support is one such example. In this situation, you can have more than one cache with the same name. The output from the **listAllCaches** option has a feature column which lists the feature that

created the cache, usually default. For multitenancy support, the feature is mt, and the cache is listed in the Incompatible shared caches section of the output.

mprotect=[all | default | none]

By default, the memory pages that contain the cache are always protected, unless a specific page is being updated. This protection helps prevent accidental or deliberate corruption to the cache. The cache header is not protected by default because this protection has a small performance cost. Specifying all ensures that all the cache pages are protected, including the header. Specifying none disables the page protection.

modified=<modified context>

Used when a JVMTI agent is installed that might modify bytecode at run time. If you do not specify this suboption and a bytecode modification agent is installed, classes are safely shared with an extra performance cost. The <modified context> is a descriptor chosen by the user; for example, *myModification1*. This option partitions the cache, so that only JVMs using context *myModification1* can share the same classes. For instance, if you run an application with a modification context and then run it again with a different modification context, all classes are stored twice in the cache. See Dealing with runtime bytecode modification for more information.

name=<name>

Connects to a cache of a given name, creating the cache if it does not exist. This option is also used to indicate the cache that is to be modified by cache utilities; for example, **destroy**. Use the **listAllCaches** utility to show which named caches are currently available. If you do not specify a name, the default name "sharedcc_%u" is used. "%u" in the cache name inserts the current user name. You can specify "%g" in the cache name to insert the current group name.

Note: Some features, when enabled, result in the creation of caches that cannot be shared with caches that were created when the feature was disabled. The multitenancy support is one such example. In this situation, you can have more than one cache with the same name.

noaot

Disables caching and loading of AOT code.

noBootclasspath

Disables the storage of classes loaded by the bootstrap class loader in the shared classes cache. Often used with the SharedClassURLFilter API to control exactly which classes are cached. See Using the SharedClassHelper API for more information about shared class filtering.

none

Added to the end of a command line, disables class data sharing. This suboption overrides class sharing arguments found earlier on the command line.

nonfatal

Allows the JVM to start even if class data sharing fails. Normal behavior for the JVM is to refuse to start if class data sharing fails. If you select **nonfatal** and the shared classes cache fails to initialize, the JVM attempts to connect to the cache in read-only mode. If this attempt fails, the JVM starts without class data sharing.

printAllStats (Utility option)

Displays detailed information about the contents of the cache that is

specified in the **name=<name>** suboption. If the name is not specified, statistics are displayed about the default cache. Every class is listed in chronological order with a reference to the location from which it was loaded. See `printAllStats` utility for more information.

printStats (Utility option)

Displays summary information for the cache that is specified by the **name**, **cacheDir**, and **nonpersistent** suboptions. The most useful information that is displayed is how full the cache is and how many classes it contains. Stale classes are classes that are updated on the file system and which the cache has therefore marked as "stale". Stale classes are not purged from the cache and can be reused. See `printStats` utility for more information.

readonly

Opens an existing cache with read-only permissions. The JVM does not create a new cache with this suboption. Opening a cache read-only prevents the JVM from making any updates to the cache. It also allows the JVM to connect to caches created by other users or groups without requiring write access. By default, this suboption is not specified.

reset

Causes a cache to be destroyed and then re-created when the JVM starts up. This option can be added to the end of a command line as **-XshareClasses:reset**.

safemode

Forces the JVM to load all classes from disk and apply the modifications to those classes (if applicable). For more information, see `Using the safemode option`.

This sub option is deprecated in IBM SDK for z/OS, V6.

silent

Disables all shared class messages, including error messages. Unrecoverable error messages, which prevent the JVM from initializing, are displayed.

verbose

Gives detailed output on the cache I/O activity, listing information about classes that are stored and found. Each class loader is given a unique ID (the bootstrap loader is always 0) and the output shows the class loader hierarchy at work, where class loaders must ask their parents for a class before they can load it themselves. It is typical to see many failed requests; this behavior is expected for the class loader hierarchy. The standard option **-verbose:class** also enables class sharing verbose output if class sharing is enabled.

verboseAOT

Enables verbose output when compiled AOT code is being found or stored in the cache. AOT code is generated heuristically. You might not see any AOT code that is generated at all for a small application. You can disable AOT caching using the **noaot** suboption. See the `IBM JVM Messages Guide` for a list of the messages produced.

verboseHelper

Enables verbose output for the Java Helper API. This output shows you how the Helper API is used by your class loader.

verboseIO

Gives detailed output on the cache I/O activity, listing information about classes that are stored and found. Each class loader is given a unique ID

(the bootstrap loader is always 0) and the output shows the class loader hierarchy at work, where class loaders must ask their parents for a class before they can load it themselves. It is typical to see many failed requests; this behavior is expected for the class loader hierarchy.

-Xsigcatch

Enables VM signal handling code.

-Xsigcatch

See also **-Xnosigcatch**. By default, signal handling is enabled.

-Xsigchain

Enables signal handler chaining.

-Xsigchain

See also **-Xnosigchain**. By default, signal handler chaining is enabled.

-Xsignal:posixSignalHandler=cooperativeShutdown

This option affects the behavior of JVM signal handlers.

-Xsignal:posixSignalHandler=cooperativeShutdown

When the JVM signal handlers for SIGSEGV, SIGILL, SIGBUS, SIGFPE, SIGTRAP, and SIGABRT end a process, they call `exit()`, by default. In this case, the z/OS Language Environment is not aware that the JVM ended abnormally.

With **-Xsignal:posixSignalHandler=cooperativeShutdown**, the JVM no longer uses `exit()` to end the process from the signal handlers. Instead, the JVM behaves in one of the following ways:

- In response to a z/OS hardware exception, the JVM uses `return()`.
- In response to signals raised or injected by software, the JVM ends the enclave with `abend 3565`.

Language Environment detects that the JVM is ending abnormally and initiates Resource Recovery Services. For more information about signal handlers, see “Signals used by the JVM” on page 39.

-Xss

Sets the maximum stack size for Java threads.

-Xss<size>

The default is 256 KB for 32-bit JVMs and 512 KB for 64-bit JVMs. The maximum value varies according to platform and specific machine configuration. If you exceed the maximum value, a `java/lang/OutOfMemoryError` message is reported.

-Xssi

Sets the stack size increment for Java threads.

-Xssi<size>

When the stack for a Java thread becomes full it is increased in size by this value until the maximum size (**-Xss**) is reached. The default is 16 KB.

-Xthr

-Xthr:<suboptions>

-Xthr:minimizeUserCPU

Minimizes user-mode CPU usage in thread synchronization where possible. The reduction in CPU usage might be a trade-off in exchange for decreased performance.

-Xtrace

Trace options.

-Xtrace[:help] | [[:<option>=<value>, ...]
See Controlling the trace for more information.

-Xtune:virtualized

Optimizes JVM function for virtualized environments, such as a cloud.

-Xtune:virtualized

Optimizes JVM function for virtualized environments, such as a cloud.

-Xverify

Use this option to enable or disable the verifier.

-Xverify[:<option>]

With no parameters, enables the verifier, which is the default. Therefore, if used on its own with no parameters, for example, **-Xverify**, this option does nothing. Optional parameters are as follows:

- **all** - enable maximum verification
- **none** - disable the verifier
- **remote** - enables strict class-loading checks on remotely loaded classes

The verifier is on by default and must be enabled for all production servers. Running with the verifier off is not a supported configuration. If you encounter problems and the verifier was turned off using **-Xverify:none**, remove this option and try to reproduce the problem.

-Xzero

Enables reduction of the memory footprint of the Java runtime environment when concurrently running multiple Java invocations.

-Xzero[:<option>]

-Xzero might not be appropriate for all types of applications because it changes the implementation of `java.util.ZipFile`, which might cause extra memory usage. **-Xzero** includes the optional parameters:

- **j9zip** - enables the j9zip sub option
- **noj9zip** - disables the j9zip sub option
- **sharezip** - enables the sharezip sub option
- **nosharezip** - disables the sharezip sub option
- **-** - enables the sharebootzip sub option
- **-** - disables the sharebootzip sub option
- **none** - disables all sub options
- **describe** - prints the sub options in effect

Because future versions might include more default options, **-Xzero** options are used to specify the sub options that you want to disable. By default, **-Xzero** enables **j9zip** and **sharezip**. A combination of **j9zip** and **sharezip** enables all .jar files to have shared caches:

- **j9zip** - uses a new `java.util.ZipFile` implementation. This suboption is not a requirement for **sharezip**; however, if **j9zip** is not enabled, only the bootstrap .jar files have shared caches.
- **sharezip** - puts the j9zip cache into shared memory. The j9zip cache is a map of zip entry names to file positions used to quickly find entries in the .zip file. You must enable **-Xshareclasses** to avoid a warning message. When using the **sharezip** suboption, note that this suboption allows every opened

.zip file and .jar file to store the j9zip cache in shared memory, so you might fill the shared memory when opening multiple new .zip files and .jar files. The affected API is `java.util.zip.ZipFile` (superclass of `java.util.jar.JarFile`). The .zip and .jar files do not have to be on a class path.

The system property `com.ibm.zero.version` is defined, and has a current value of 2. Although **-Xzero** is accepted on all platforms, support for the sub options varies by platform:

- **-Xzero** with all other sub options are available only on Windows x86-32 and Linux x86-32 platforms.

JVM -XX command-line options

JVM command-line options that are specified with -XX are not recommended for casual use.

These options are subject to change without notice.

-XXallowvmshutdown

This option is provided as a workaround for customer applications that cannot shut down cleanly, as described in APAR IZ59734.

-XXallowvmshutdown: [false|true]

Customers who need this workaround should use **-XXallowvmshutdown: false**. The default option is **-XXallowvmshutdown: true** for Java 6 SR5 onwards.

-XX:codectotal

Use this option to set the maximum size limit for the JIT code cache.

-XX:codectotal=<size>

This option is an alias for the “-Xcodectotal” on page 99 option.

-XX:MaxDirectMemorySize

Sets the maximum size for an nio direct buffer.

-XX:MaxDirectMemorySize=<size>

When you set a value for this property, the size cannot exceed that value. If you do not set a value, a soft limit of 64 MB is set. The JVM automatically expands this soft limit in 32 MB chunks, as required.

-XX:-StackTraceInThrowable

This option removes stack traces from exceptions.

-XX:-StackTraceInThrowable

By default, stack traces are available in exceptions. Including a stack trace in exceptions requires walking the stack and that can affect performance. Removing stack traces from exceptions can improve performance but can also make problems harder to debug.

When this option is enabled, `Throwable.getStackTrace()` returns an empty array and the stack trace is displayed when an uncaught exception occurs. `Thread.getStackTrace()` and `Thread.getAllStackTraces()` are not affected by this option.

-XX:[+|-]UseCompressedOops (64-bit only)

This option enables or disables compressed references in 64-bit JVMs, and is provided to help when porting applications from the Oracle JVM to the IBM JVM. This option might not be supported in subsequent releases.

-XX: [+ | -] UseCompressedOops

The **-XX:+UseCompressedOops** option enables compressed references in 64-bit JVMs. The **-XX:+UseCompressedOops** option is similar to specifying **-Xcompressedrefs**, which is detailed in the topic “JVM command-line options” on page 82.

The **-XX:-UseCompressedOops** option prevents the use of compressed references in 64-bit JVMs.

JIT and AOT command-line options

Use these JIT and AOT compiler command-line options to control code compilation.

For options that take a *<size>* parameter, suffix the number with “k” or “K” to indicate kilobytes, “m” or “M” to indicate megabytes, or “g” or “G” to indicate gigabytes.

For more information about JIT and AOT, see JIT and AOT problem determination.

-Xaot

Use this option to control the behavior of the AOT compiler.

-Xaot[:<parameter>=<value>, ...]

With no parameters, enables the AOT compiler. The AOT compiler is enabled by default but is not active unless shared classes are enabled. Using this option on its own has no effect. The following parameters are useful:

count=<n>

Where *<n>* is the number of times a method is called before it is compiled or loaded from an existing shared class cache. For example, setting `count=0` forces the AOT compiler to compile everything on first execution.

limitFile=(<filename>, <m>, <n>)

Compile or load only the methods listed on lines *<m>* to *<n>* in the specified limit file. Methods not listed in the limit file and methods listed on lines outside the range are not compiled or loaded.

loadExclude=<methods>

Do not load methods beginning with *<methods>*.

loadLimit=<methods>

Load methods beginning with *<methods>* only.

loadLimitFile=(<filename>, <m>, <n>)

Load only the methods listed on lines *<m>* to *<n>* in the specified limit file. Methods not listed in the limit file and methods listed on lines outside the range are not loaded.

verbose

Reports information about the AOT and JIT compiler configuration and method compilation.

-Xcodecache

This option is used to tune performance.

-Xcodecache<size>

This option sets the size of each block of memory that is allocated to store the native code of compiled Java methods. By default, this size is selected internally according to the processor architecture and the capability of your system. If profiling tools show significant costs in *trampolines*, that is a good reason to change the size until the costs are reduced. Changing the size does not mean always increasing the size. The option provides the mechanism to tune for the correct size until hot interblock calls are eliminated. A reasonable starting point to tune for the optimal size is (totalNumberByteOfCompiledMethods * 1.1).

Note: Trampolines are small pieces of code that facilitate the transition between two compiled methods that are far apart from each other in memory. Jvmti identifies trampolines in a methodLoad2 event.

-Xcodecachetotal

Use this option to set the maximum size limit for the JIT code cache.

-Xcodecachetotal<size>

See “JIT and AOT command-line options” on page 98 for more information about the <size> parameter.

By default, the total size of the JIT code cache is determined by your operating system, architecture, and the version of IBM SDK Java Technology Edition that you are using. Long-running, complex, server-type applications can fill the JIT code cache, which can cause performance problems because not all of the important methods can be JIT-compiled. Use the **-Xcodecachetotal** option to increase the maximum code cache size beyond the default setting, to a setting that suits your application.

The value that you specify is rounded up to a multiple of the code cache block size, as specified by the “-Xcodecache” on page 98 option. If you specify a value for the **-Xcodecachetotal** option that is smaller than the default setting, that value is ignored.

The maximum size limits, for both the JIT code and data caches, that are in use by the JVM are shown in Jvdump output. Look for lines that begin with 1STSEGLIMIT. Use this information together with verbose JIT tracing to determine suitable values for this option on your system. For example Jvdump output, see Storage Management (MEMINFO).

Related reference:

“-Xjit” on page 100

Use the JIT compiler command line option to produce verbose JIT trace output.

Related information:

Using Jvdump

-Xcomp (z/OS only)

Forces methods to be compiled by the JIT compiler on their first use.

-Xcomp

The use of this option is deprecated; use **-Xjit:count=0** instead.

-Xint

This option makes the JVM use the Interpreter only, disabling the Just-In-Time (JIT) and Ahead-Of-Time (AOT) compilers.

-Xint

By default, the JIT compiler is enabled. By default, the AOT compiler is enabled, but is not used by the JVM unless shared classes are also enabled.

-Xjit

Use this option to control the behavior of the JIT compiler.

-Xjit[:<parameter>=<value>, ...]

With no parameters, enables the JIT compiler. The JIT compiler is enabled by default, so using this option on its own has no effect. Useful parameters are:

count=<n>

Where <n> is the number of times a method is called before it is compiled. For example, setting count=0 forces the JIT compiler to compile everything on first execution.

exclude={<method>}

Excludes the specified method from compilation.

limitFile=(<filename>, <m>, <n>)

Compile only the methods listed on lines <m> to <n> in the specified limit file. Methods not listed in the limit file and methods listed on lines outside the range are not compiled.

optlevel=[noOpt | cold | warm | hot | veryHot | scorching]

Forces the JIT compiler to compile all methods at a specific optimization level. Specifying **optlevel** might have an unexpected effect on performance, including reduced overall performance.

verbose[={compileStart|compileEnd}]

Reports information about the JIT and AOT compiler configuration and method compilation.

The **={compileStart|compileEnd}** option reports when the JIT starts to compile a method, and when it ends.

vlog=<filename>

Sends verbose output to a file. If you do not specify this parameter, the output is sent to the standard error output stream (stderr).

Related information:

Diagnosing a JIT or AOT problem

-Xnoaot

This option turns off the AOT compiler and disables the use of AOT-compiled code.

-Xnoaot

By default, the AOT compiler is enabled but is active only when shared classes are also enabled. Using this option does not affect the JIT compiler.

-Xnojit

This option turns off the JIT compiler.

-Xnojit

By default, the JIT compiler is enabled. This option does not affect the AOT compiler.

-Xquickstart

This option causes the JIT compiler to run with a subset of optimizations.

-Xquickstart

The effect is faster compilation times that improve startup time, but longer running applications might run slower. When the AOT compiler is active (both shared classes and AOT compilation enabled), **-Xquickstart** causes all methods to be AOT compiled. The AOT compilation improves the startup time of subsequent runs, but might reduce performance for longer running applications. **-Xquickstart** can degrade performance if it is used with long-running applications that contain hot methods. The implementation of **-Xquickstart** is subject to change in future releases. By default, **-Xquickstart** is disabled..

-XsamplingExpirationTime

Use this option to disable JIT sampling after a specified amount of time.

-XsamplingExpirationTime<time>

Disables the JIT sampling thread after <time> seconds. When the JIT sampling thread is disabled, no processor cycles are used by an idle JVM.

-Xscmaxaot

Optionally applies a maximum number of bytes in the class cache that can be used for AOT data.

-Xscmaxaot<size>

This option is useful if you want a certain amount of cache space guaranteed for non-AOT data. If this option is not specified, the maximum limit for AOT data is the amount of free space in the cache. The value of this option must not be smaller than the value of **-Xscminaot** and must not be larger than the value of **-Xscmx**.

-Xscminaot

Optionally applies a minimum number of bytes in the class cache to reserve for AOT data.

-Xscminaot<size>

If this option is not specified, no space is reserved for AOT data. However, AOT data is still written to the cache until the cache is full or the **-Xscmaxaot** limit is reached. The value of this option must not exceed the value of **-Xscmx** or **-Xscmaxaot**. The value of **-Xscminaot** must always be considerably less than the total cache size, because AOT data can be created only for cached classes. If the value of **-Xscminaot** equals the value of **-Xscmx**, no class data or AOT data can be stored.

Garbage Collector command-line options

Use these Garbage Collector command-line options to control garbage collection.

You might need to read Memory management to understand some of the references that are given here.

The **-verbose:gc** option detailed in Verbose garbage collection logging is the main diagnostic aid that is available for runtime analysis of the Garbage Collector. However, additional command-line options are available that affect the behavior of the Garbage Collector and might aid diagnostic data collection.

For options that take a <size> parameter, suffix the number with "k" or "K" to indicate kilobytes, "m" or "M" to indicate megabytes, or "g" or "G" to indicate gigabytes.

For options that take a *<percentage>* parameter, use a number from 0 to 1, for example, 50% is 0.5.

-Xalwaysclassgc

Always perform dynamic class unloading checks during global collection.

-Xalwaysclassgc

The default behavior is as defined by **-Xclassgc**.

-Xclassgc

Enables dynamic unloading of classes by the JVM. Garbage collection of class objects occurs only on class loader changes.

-Xclassgc

Dynamic unloading is the default behavior. To disable dynamic class unloading, use the **-Xnoclassgc** option.

-Xcompactexplicitgc

Enables full compaction each time System.gc() is called.

-Xcompactexplicitgc

Enables full compaction each time System.gc() is called.

-Xcompactgc

Compacts on all garbage collections (system and global).

-Xcompactgc

The default (no compaction option specified) makes the GC compact based on a series of triggers that attempt to compact only when it is beneficial to the future performance of the JVM.

-Xconcurrentbackground

Specifies the number of low-priority background threads attached to assist the mutator threads in concurrent mark.

-Xconcurrentbackground<number>

The default is 0 on Linux System z and 1 on all other platforms.

-Xconcurrentlevel

Specifies the allocation "tax" rate.

-Xconcurrentlevel<number>

This option indicates the ratio between the amount of heap allocated and the amount of heap marked. The default is 8.

-Xconcurrentslack

Attempts to keep the specified amount of the heap space free in concurrent collectors by starting the concurrent operations earlier.

-Xconcurrentslack<size>

Using this option can sometimes alleviate pause time problems in concurrent collectors at the cost of longer concurrent cycles, affecting total throughput. The default value is 0, which is optimal for most applications.

-Xconmeter

This option determines the usage of which area, LOA (Large Object Area) or SOA (Small Object Area), is metered and hence which allocations are taxed during concurrent mark.

-Xconmeter:<soa | loa | dynamic>

Using **-Xconmeter:soa** (the default) applies the allocation tax to allocations

from the small object area (SOA). Using **-Xconmeter:loa** applies the allocation tax to allocations from the large object area (LOA). If **-Xconmeter:dynamic** is specified, the collector dynamically determines which area to meter based on which area is exhausted first, whether it is the SOA or the LOA.

-Xdisableexcessivegc

Disables the throwing of an OutOfMemory exception if excessive time is spent in the GC.

-Xdisableexcessivegc

Disables the throwing of an OutOfMemory exception if excessive time is spent in the GC.

-Xdisableexplicitgc

Disables System.gc() calls.

-Xdisableexplicitgc

Many applications still make an excessive number of explicit calls to System.gc() to request garbage collection. In many cases, these calls degrade performance through premature garbage collection and compactions. However, you cannot always remove the calls from the application.

The **-Xdisableexplicitgc** parameter allows the JVM to ignore these garbage collection suggestions. Typically, system administrators use this parameter in applications that show some benefit from its use.

By default, calls to System.gc() trigger a garbage collection.

-Xdisablestringconstantgc

Prevents strings in the string intern table from being collected.

-Xdisablestringconstantgc

Prevents strings in the string intern table from being collected.

-Xenableexcessivegc

If excessive time is spent in the GC, the option returns null for an allocate request and thus causes an OutOfMemory exception to be thrown.

-Xenableexcessivegc

The OutOfMemory exception is thrown only when the heap has been fully expanded and the time spent is making up at least 95%. This behavior is the default.

You can control the percentage that triggers an excessive GC event with the **-Xgc:excessiveGCratio** option. For more information, see “-Xgc.”

-Xenablestringconstantgc

Enables strings from the string intern table to be collected.

-Xenablestringconstantgc

This option is on by default.

-Xgc

Options that change the behavior of the Garbage Collector (GC). These options are deprecated.

-Xgc:<excessiveGCratio | verbose | compact | nocompact >

excessiveGCratio=*value*

Where *value* is a percentage. The default value is 95. This option can be used only when **-Xenableexcessivegc** is set. For more information, see “-Xenableexcessivegc” on page 103.

Options **verbose**, **compact**, and **nocompact** are deprecated.

-Xgcpolicy

Controls the behavior of the Garbage Collector.

-Xgcpolicy:< gencon | optavgpause | optthruput | subpool (AIX, Linux and IBM i on IBM POWER® architecture, Linux and z/OS on zSeries) >

gencon

The **gencon** policy uses a concurrent mark phase combined with generational garbage collection to help minimize the time that is spent in any garbage collection pause. This policy is particularly useful for applications with many short-lived objects, such as transactional applications. Pause times can be significantly shorter than with the **optthruput** policy, while still producing good throughput. Heap fragmentation is also reduced.

optavgpause

The **optavgpause** policy uses concurrent mark and concurrent sweep phases. Pause times are shorter than with **optthruput**, but application throughput is reduced because some garbage collection work is taking place while the application is running. Consider using this policy if you have a large heap size (available on 64-bit platforms), because this policy limits the effect of increasing heap size on the length of the garbage collection pause. However, if your application uses many short-lived objects, the **gencon** policy might produce better performance.

subpool

The **subpool** policy disables the concurrent mark phase, but employs an improved object allocation algorithm to achieve better performance when allocating objects on the heap. This algorithm is more suitable for multiple processor systems, commonly 16 processors or more. Applications that must scale on large systems might benefit from this policy. This policy is available on AIX, Linux PPC and zSeries, z/OS, and i5/OS only.

optthruput

The **optthruput** policy (default) disables the concurrent mark phase. The application stops during global garbage collection, so long pauses can occur. This configuration is typically used for large-heap applications when high application throughput, rather than short garbage collection pauses, is the main performance goal. If your application cannot tolerate long garbage collection pauses, consider using another policy, such as **gencon**.

-Xgcthreads

Sets the number of threads that the Garbage Collector uses for parallel operations.

-Xgcthreads<number>

The total number of GC threads is composed of one application thread with the remainder being dedicated GC threads. By default, the number is set to $n-1$, where n is the number of reported CPUs. Where SMT or hyperthreading is in place, the number of reported CPUs is larger than the number of physical CPUs. Likewise, where virtualization is in place, the number of reported CPUs is the number of virtual CPUs assigned to the operating system. To set it to a different number, for example 4, use **-Xgcthreads4**. The minimum valid value

is 1, which disables parallel operations, at the cost of performance. No advantage is gained if you increase the number of threads to more than the default setting.

On systems running multiple JVMs or in LPAR environments where multiple JVMs can share the same physical CPUs, you might want to restrict the number of GC threads used by each JVM. The restriction helps prevent the total number of parallel operation GC threads for all JVMs exceeding the number of physical CPUs present, when multiple JVMs perform garbage collection at the same time.

-Xgcworkpackets

Specifies the total number of work packets available in the global collector.

-Xgcworkpackets<number>

If you do not specify a value, the collector allocates a number of packets based on the maximum heap size.

-Xloa

Allocates a large object area (LOA).

-Xloa

Objects are allocated in this LOA rather than the SOA. By default, the LOA is enabled for all GC policies except for subpool, where the LOA is not available.

-Xloainitial

Specifies the initial percentage (between 0 and 0.95) of the current tenure space allocated to the large object area (LOA).

-Xloainitial<percentage>

The default value is 0.05, which is 5%.

-Xloamaximum

Specifies the maximum percentage (between 0 and 0.95) of the current tenure space allocated to the large object area (LOA).

-Xloamaximum<percentage>

The default value is 0.5, which is 50%.

-Xloaminimum

Specifies the minimum percentage (between 0 and 0.95) of the current tenure space allocated to the large object area (LOA).

-Xloaminimum<percentage>

The LOA does not shrink to less than this value. The default value is 0, which is 0%.

-Xmaxe

Sets the maximum amount by which the garbage collector expands the heap.

-Xmaxe<size>

Typically, the garbage collector expands the heap when the amount of free space falls to less than 30% (or by the amount specified using **-Xminf**), by the amount required to restore the free space to 30%. The **-Xmaxe** option limits the expansion to the specified value; for example **-Xmaxe10M** limits the expansion to 10 MB. By default, there is no maximum expansion size.

-Xmaxf

Specifies the maximum percentage of heap that must be free after a garbage collection.

-Xmaxf<percentage>

If the free space exceeds this amount, the JVM tries to shrink the heap. The default value is 0.6 (60%).

-Xmaxt

Specifies the maximum percentage of time to be spent in Garbage Collection.

-Xmaxt<percentage>

If the percentage of time exceeds this value, the JVM tries to expand the heap. The default value is 13%.

-Xmca

Sets the expansion step for the memory allocated to store the RAM portion of loaded classes.

-Xmca<size>

Each time more memory is required to store classes in RAM, the allocated memory is increased by this amount. By default, the expansion step is 32 KB. Use the **-verbose:sizes** option to determine the value that the VM is using. If the expansion step size you choose is too large, `OutOfMemoryError` is reported. The exact value of a “too large” expansion step size varies according to the platform and the specific machine configuration.

-Xmco

Sets the expansion step for the memory allocated to store the ROM portion of loaded classes.

-Xmco<size>

Each time more memory is required to store classes in ROM, the allocated memory is increased by this amount. By default, the expansion step is 128 KB. Use the **-verbose:sizes** option to determine the value that the VM is using. If the expansion step size you choose is too large, `OutOfMemoryError` is reported. The exact value of a “too large” expansion step size varies according to the platform and the specific machine configuration.

-Xmine

Sets the minimum amount by which the Garbage Collector expands the heap.

-Xmine<size>

Typically, the garbage collector expands the heap by the amount required to restore the free space to 30% (or the amount specified using **-Xminf**). The **-Xmine** option sets the expansion to be at least the specified value; for example, **-Xmine50M** sets the expansion size to a minimum of 50 MB. By default, the minimum expansion size is 1 MB.

-Xminf

Specifies the minimum percentage of heap to remain free after a garbage collection.

-Xminf<percentage>

If the free space falls to less than this amount, the JVM attempts to expand the heap. The default value is 30%.

-Xmint

Specifies the minimum percentage of time to spend in Garbage Collection.

-Xmint<percentage>

If the percentage of time drops to less than this value, the JVM tries to shrink the heap. The default value is 5%.

-Xmn

Sets the initial and maximum size of the new area to the specified value when using **-Xgcpolicy:gencon**.

-Xmn<size>

Equivalent to setting both **-Xmns** and **-Xmnx**. If you set either **-Xmns** or **-Xmnx**, you cannot set **-Xmn**. If you try to set **-Xmn** with either **-Xmns** or **-Xmnx**, the VM does not start, returning an error. By default, **-Xmn** is not set. If the scavenger is disabled, this option is ignored.

-Xmns

Sets the initial size of the new area to the specified value when using **-Xgcpolicy:gencon**.

-Xmns<size>

By default, this option is set to 25% of the value of the **-Xms** option. This option returns an error if you try to use it with **-Xmn**. You can use the **-verbose:sizes** option to find out the values that the VM is currently using. If the scavenger is disabled, this option is ignored.

-Xmnx

Sets the maximum size of the new area to the specified value when using **-Xgcpolicy:gencon**.

-Xmnx<size>

By default, this option is set to 25% of the value of the **-Xmx** option. This option returns an error if you try to use it with **-Xmn**. You can use the **-verbose:sizes** option to find out the values that the VM is currently using. If the scavenger is disabled, this option is ignored.

-Xmo

Sets the initial and maximum size of the old (tenured) heap to the specified value when using **-Xgcpolicy:gencon**.

-Xmo<size>

Equivalent to setting both **-Xmos** and **-Xmox**. If you set either **-Xmos** or **-Xmox**, you cannot set **-Xmo**. If you try to set **-Xmo** with either **-Xmos** or **-Xmox**, the VM does not start, returning an error. By default, **-Xmo** is not set.

-Xmoi

Sets the amount the Java heap is incremented when using **-Xgcpolicy:gencon**.

-Xmoi<size>

If set to zero, no expansion is allowed. By default, the increment size is calculated on the expansion size, set by **-Xmine** and **-Xminf**.

-Xmos

Sets the initial size of the old (tenure) heap to the specified value when using **-Xgcpolicy:gencon**.

-Xmos<size>

By default, this option is set to 75% of the value of the **-Xms** option. This option returns an error if you try to use it with **-Xmo**. You can use the **-verbose:sizes** option to find out the values that the VM is currently using.

-Xmox

Sets the maximum size of the old (tenure) heap to the specified value when using **-Xgcpolicy:gencon**.

-Xmx<size>

By default, this option is set to the same value as the **-Xmx** option. This option returns an error if you try to use it with **-Xmo**. You can use the **-verbose:sizes** option to find out the values that the VM is currently using.

-Xmr

Sets the size of the Garbage Collection "remembered set".

-Xmr<size>

The Garbage Collection "remembered set" is a list of objects in the old (tenured) heap that have references to objects in the new area. By default, this option is set to 16 K.

-Xmrx

Sets the remembered maximum size setting.

-Xmrx<size>

Sets the remembered maximum size setting.

-Xms

Sets the initial Java heap size.

-Xms<size>

<size> can be specified in megabytes (m) or gigabytes (g). For example: **-Xms2g** sets an initial Java heap size of 2GB. The minimum size is 1 MB.

You can also use the **-Xmo** option.

If the scavenger is enabled, **-Xms >= -Xmn + -Xmo**.

If the scavenger is disabled, **-Xms >= -Xmo**.

Note: The **-Xmo** option is not supported by the balanced garbage collection policy.

-Xmx

Sets the maximum memory size for the application (**-Xmx >= -Xms**).

-Xmx<size>

<size> can be specified in megabytes (m) or gigabytes (g). For example: **-Xmx2g** sets a maximum heap size of 2GB.

For information about default values, see "Default settings for the JVM" on page 111.

If you are allocating the Java heap with large pages, read the information provided for the **"-Xlp"** on page 88 option.

Examples of the use of **-Xms** and **-Xmx**:

-Xms2m -Xmx64m

Heap starts at 2 MB and grows to a maximum of 64 MB.

-Xms100m -Xmx100m

Heap starts at 100 MB and never grows.

-Xms20m -Xmx1024m

Heap starts at 20 MB and grows to a maximum of 1 GB.

-Xms50m

Heap starts at 50 MB and grows to the default maximum.

-Xmx256m

Heap starts at default initial value and grows to a maximum of 256 MB.

If you exceed the limit set by the **-Xmx** option, the JVM generates an `OutOfMemoryError`.

-Xnoclassgc

Disables class garbage collection.

-Xnoclassgc

This option switches off garbage collection of storage associated with Java classes that are no longer being used by the JVM. The default behavior is as defined by **-Xclassgc**. Enabling this option is not recommended except under the direction of the IBM Java support team. The reason is the option can cause unlimited native memory growth, leading to out-of-memory errors.

-Xnocompactexplicitgc

Disables compaction on `System.gc()` calls.

-Xnocompactexplicitgc

Compaction takes place on global garbage collections if you specify **-Xcompactgc** or if compaction triggers are met. By default, compaction is enabled on calls to `System.gc()`.

-Xnocompactgc

Disables compaction on all garbage collections (system or global).

-Xnocompactgc

By default, compaction is enabled.

-Xnoloa

Prevents allocation of a large object area; all objects are allocated in the SOA.

-Xnoloa

See also **-Xloa**.

-Xnopartialcompactgc

Disables incremental compaction.

-Xnopartialcompactgc

See also **-Xpartialcompactgc**.

-Xpartialcompactgc

Enables incremental compaction.

-Xpartialcompactgc

See also **-Xnopartialcompactgc**. By default, this option is not set, so all compactations are full.

-Xsoftmx

This option sets a "soft" maximum limit for the initial size of the Java heap.

-Xsoftmx<size> (AIX only)

Use the **-Xmx** option to set a "hard" limit for the maximum size of the heap. By default, **-Xsoftmx** is set to the same value as **-Xmx**. The value of **-Xms** must be less than, or equal to, the value of **-Xsoftmx**. See the introduction to this topic for more information about specifying `<size>` parameters.

You can set this option on the command line, then modify it at run time by using the `MemoryMXBean.setMaxHeapSize()` method in the `com.ibm.lang.management` API. By using this API, Java applications can

dynamically monitor and adjust the heap size as required. This function can be useful in virtualized or cloud environments, for example, where the available memory might change dynamically to meet business needs. When you use the API, you must specify the value in bytes, such as 2147483648 instead of 2g.

For example, you might set the initial heap size to 1 GB and the maximum heap size to 8 GB. You might set a smaller value, such as 2 GB, for **-Xsoftmx**, to limit the heap size that is used initially:

```
-Xms1g -Xsoftmx2g -Xmx8g
```

You can then use the `com.ibm.lang.management` API from within a Java application to increase the **-Xsoftmx** value during run time, as load increases. This change allows the application to use more memory than you specified initially.

-Xsoftrefthreshold

Sets the value used by the garbage collector to determine the number of garbage collections after which a soft reference is cleared if its referent has not been marked.

-Xsoftrefthreshold<number>

The default is 32, meaning that the soft reference is cleared after 32 * (percentage of free heap space) garbage collection cycles where its referent was not marked. For example, if **-Xsoftrefthreshold** is set to 32, and the heap is 50% free, soft references are cleared after 16 garbage collection cycles.

-Xtgc

Provides garbage collection tracing options.

-Xtgc:<arguments>

<arguments> is a comma-separated list containing one or more of the following arguments:

backtrace

Before a garbage collection, a single line is printed containing the name of the master thread for garbage collection, as well as the value of the `osThread` slot in the `J9VMThread` structure.

compaction

Prints extra information showing the relative time spent by threads in the "move" and "fixup" phases of compaction

concurrent

Prints extra information showing the activity of the concurrent mark background thread

dump

Prints a line of output for every free chunk of memory in the system, including "dark matter" (free chunks that are not on the free list for some reason, typically because they are too small). Each line contains the base address and the size in bytes of the chunk. If the chunk is followed in the heap by an object, the size and class name of the object is also printed. This argument has a similar effect to the **terse** argument.

freeList

Before a garbage collection, prints information about the free list and allocation statistics since the last garbage collection. Prints the number of items on the free list, including "deferred" entries (with the scavenger, the unused space is a deferred free list entry). For TLH and non-TLH allocations, prints the total number of allocations, the average allocation

size, and the total number of bytes discarded during allocation. For non-TLH allocations, also included is the average number of entries that were searched before a sufficiently large entry was found.

parallel

Produces statistics on the activity of the parallel threads during the mark and sweep phases of a global garbage collection.

references

Prints extra information every time that a reference object is enqueued for finalization, showing the reference type, reference address, and referent address.

scavenger

Prints extra information after each scavenger collection. A histogram is produced showing the number of instances of each class, and their relative ages, present in the survivor space. The information is obtained by performing a linear walk-through of the space.

terse

Dumps the contents of the entire heap before and after a garbage collection. For each object or free chunk in the heap, a line of trace output is produced. Each line contains the base address, "a" if it is an allocated object, and "f" if it is a free chunk, the size of the chunk in bytes, and, if it is an object, its class name.

-Xverbosegclog

Causes **-verbose:gc** output to be written to a specified file.

-Xverbosegclog[:<file>[,<X>,<Y>]]

If the file cannot be found, **-verbose:gc** tries to create the file, and then continues as normal if it is successful. If it cannot create the file (for example, if an invalid filename is passed into the command), it redirects the output to stderr.

If you specify **<X>** and **<Y>** the **-verbose:gc** output is redirected to **X** files, each containing **Y** GC cycles.

The dump agent tokens can be used in the filename. See Dump agent tokens for more information. If you do not specify **<file>**, **verbosegc.%Y%m%d.%H%M%S.%pid.txt** is used.

By default, no verbose GC logging occurs.

Default settings for the JVM

This appendix shows the default settings that the JVM uses. These settings affect how the JVM operates if you do not apply any changes to its environment. The tables show the JVM operation and the default setting.

These tables are a quick reference to the state of the JVM when it is first installed. The last column shows how the default setting can be changed:

- c** The setting is controlled by a command-line parameter only.
- e** The setting is controlled by an environment variable only.
- ec** The setting is controlled by a command-line parameter or an environment variable. The command-line parameter always takes precedence.

JVM setting	Default	Setting affected by
Javadump	Enabled	ec
Heapdump	Disabled	ec
System dump	Enabled	ec
Snap traces	Enabled	ec
Verbose output	Disabled	c
Boot classpath search	Disabled	c
JNI checks	Disabled	c
Remote debugging	Disabled	c
Strict conformance checks	Disabled	c
Quickstart	Disabled	c
Remote debug info server	Disabled	c
Reduced signaling	Disabled	c
Signal handler chaining	Enabled	c
Classpath	Not set	ec
Class data sharing	Disabled	c
Accessibility support	Enabled	e
JIT compiler	Enabled	ec
AOT compiler (AOT is not used by the JVM unless shared classes are also enabled)	Enabled	c
JIT debug options	Disabled	c
Java2D max size of fonts with algorithmic bold	14 point	e
Java2D use rendered bitmaps in scalable fonts	Enabled	e
Java2D freetype font rasterizing	Enabled	e
Java2D use AWT fonts	Disabled	e

JVM setting	AIX	IBM i	Linux	Windows	z/OS	Setting affected by
Default locale	None	None	None	N/A	None	e
Time to wait before starting plug-in	N/A	N/A	Zero	N/A	N/A	e
Temporary directory	/tmp	/tmp	/tmp	c:\temp	/tmp	e
Plug-in redirection	None	None	None	N/A	None	e
IM switching	Disabled	Disabled	Disabled	N/A	Disabled	e
IM modifiers	Disabled	Disabled	Disabled	N/A	Disabled	e
Thread model	N/A	N/A	N/A	N/A	Native	e
Initial stack size for Java Threads 32-bit. Use: -Xiss<size>	2 KB	2 KB	2 KB	2 KB	2 KB	c
Maximum stack size for Java Threads 32-bit. Use: -Xss<size>	256 KB	256 KB	256 KB	256 KB	256 KB	c
Stack size for OS Threads 32-bit. Use -Xmso<size>	256 KB	256 KB	256 KB	32 KB	256 KB	c

JVM setting	AIX	IBM i	Linux	Windows	z/OS	Setting affected by
Initial stack size for Java Threads 64-bit. Use: -Xiss<size>	2 KB	N/A	2 KB	2 KB	2 KB	c
Maximum stack size for Java Threads 64-bit. Use: -Xss<size>	512 KB	N/A	512 KB	512 KB	512 KB	c
Stack size for OS Threads 64-bit. Use -Xmso<size>	256 KB	N/A	256 KB	256 KB	256 KB	c
Initial heap size. Use -Xms<size>	4 MB	4 MB	4 MB	4 MB	4 MB	c
Maximum Java heap size. Use -Xmx<size>	Half the available memory with a minimum of 16 MB and a maximum of 512 MB	2 GB	Half the available memory with a minimum of 16 MB and a maximum of 512 MB	Half the real memory with a minimum of 16 MB and a maximum of 2 GB	Half the available memory with a minimum of 16 MB and a maximum of 512 MB	c

“Available memory” is defined as being the smallest of two values:

- The real or “physical” memory.
- The **RLIMIT_AS** value.

Known issues and limitations

Known issues or limitations that you might encounter in specific system environments, or configurations.

If you find a problem, see the “Hints and Tips” pages, at <http://www.ibm.com/systems/z/os/zos/tools/java/faq/javafaq.html>.

If you find a problem that you have been unable to solve after looking through the “Hints and Tips” pages, see <http://www.ibm.com/systems/z/os/zos/tools/java/services/services.html> for advice and information about how to raise problems.

Limitation on class path length

If there are more than 2031 characters in your class path, the shell truncates your class path to 2031 characters. If you need a class path longer than 2031 characters, use the extension class loader option to refer to directories containing your `.jar` files, for example:

```
-Djava.ext.dirs=<directory>
```

Where `<directory>` is the directory containing your `.jar` files.

JConsole monitoring tool Local tab

In the IBM JConsole tool, the **Local** tab, which allows you to connect to other Virtual Machines on the same system, is not available. Also, the corresponding command line **pid** option is not supported. Instead, use the **Remote** tab in JConsole to connect to the Virtual Machine that you want to monitor. Alternatively,

use the **connection** command-line option, specifying a host of localhost and a port number. When you start the application that you want to monitor, set these command-line options:

-Dcom.sun.management.jmxremote.port=<value>

Specifies the port the management agent listens on.

-Dcom.sun.management.jmxremote.authenticate=false

Disables authentication unless you have created a user name file.

-Dcom.sun.management.jmxremote.ssl=false

Disables SSL encryption.

Incorrect stack traces when loading new classes after an Exception is caught

If new classes are loaded after an Exception has been caught, the stack trace contained in the Exception might become incorrect. The stack trace becomes incorrect if classes in the stack trace are unloaded, and new classes are loaded into their memory segments.

ThreadMXBean Thread User CPU Time limitation

There is no way to distinguish between user mode CPU time and system mode CPU time on this platform. `ThreadMXBean.getThreadUserTime()`, `ThreadMXBean.getThreadCpuTime()`, `ThreadMXBean.getCurrentThreadUserTime()`, and `ThreadMXBean.getCurrentThreadCpuTime()` all return the total CPU time for the required thread.

You can get the CPU time only for the current thread by calling `ThreadMXBean.isCurrentThreadCpuTimeSupported()`. Calling `ThreadMXBean.isThreadCpuTimeSupported()` returns a value of false because getting the CPU time for a thread other than the current thread is not supported.

NullPointerException with the GTK Look and Feel

DBCS environments only

If your application fails with a `NullPointerException` using the GTK Look and Feel, unset the `GNOME_DESKTOP_SESSION_ID` environment variable.

ASCII to EBCDIC

Because z/OS uses the EBCDIC character encoding instead of the more common ASCII encoding, sometimes there are portability problems with Java code written on z/OS. Inside the scope of the JVM, all character and string data is stored and manipulated in Unicode. I/O data outside of the virtual machine, such as on a disk or on a network, is converted to the native platform encoding. However, Java applications that implicitly assume ASCII in specific situations might require some alterations to run as expected under z/OS. For example, a platform-neutral application might have hard-coded dependencies, such as literals in ASCII.

The Java language contains the abstractions necessary to handle the switch between character encoding. The various Reader and Writer classes in the `java.io` package provide alternate constructors with a specified code page. This mechanism is used for globalization support, and it can also be used to force ASCII (or other)

I/O where required. Not all I/O needs to be overridden; for example, character output to the display remains in the native encoding.

In addition to the Reader and Writer classes, there are a few specific situations that might require additional care. For example, the String class has an overloaded `getBytes()` method that takes an encoding as an additional parameter. The overloaded method is useful for direct string manipulation when you are implementing custom data streams or network protocols directly in Java.

In general, straightforward workarounds are available for character encoding problems. Some encoding problems are not visible to the application because they are handled in programs running on z/OS. An example is when using Java Database Connectivity (JDBC).

IPv6 multicast support

z/OS V1R6 currently does not support IPv4-mapped Multicast addresses. If you are using an IPv4 Multicast address, you cannot join a Multicast group unless you disable IPv6 support by setting the `java.net.preferIPv4Stack` property to true.

Use the following instruction to set the property on the command line:

```
java -Djava.net.preferIPv4Stack=true <classname>
```

Unicode Shift_JIS code page alias

Japanese users only

The Unicode code page alias “`\u30b7\u30d5\u30c8\u7b26\u53f7\u5316\u8868\u73fe`” for Shift_JIS has been removed. If you use this code page in your applications, replace it with Shift_JIS.

-Xshareclasses:<options>

Shared classes cache and control files are not compatible between Java 6 SR4 and previous releases.

Java Kernel installation

The Java kernel aims to reduce the startup time imposed by an application when it finds that the installed JRE needs an update. When this situation occurs, the Java kernel automatically downloads only the Java components that are needed directly from the Oracle Web site. This automated download is currently not possible with the IBM implementation of this Oracle Java update.

Java Deployment Toolkit

The toolkit implements the JavaScript `DeployJava.js`, which can be used to automatically generate any HTML needed to deploy applets and Java Web Start applications. However, the automatic generation is not possible with the IBM release of Java, because the process involves downloading and running the specific JRE from a public site, using public functions.

Expired GTE Cybertrust Certificate

The IBM Runtime Environment for Java contains an expired GTE CyberTrust Certificate in the CACERTS file for compatibility reasons. The CACERTS file is provided as a default truststore. Some common public certificates are provided as a convenience.

If no applications require the certificate, it can be left in the CACERTS file. Alternatively, the certificate can be deleted. If applications do require the certificate, modify them to use the newer GTE CyberTrust Global root certificate that expires in 2018.

This certificate might be removed for later versions of the SDK.

Using `-Xshareclasses:destroy` during JVM startup

When running the command `java -Xshareclasses:destroy` on a shared cache that is being used by a second JVM during startup, you might have the following issues:

- The second JVM fails.
- The shared cache is deleted.

Problems accessing archive files created using `java.util.Zip*`

The `java.util.Zip*` files can create archive files that are larger than 4 GB. However, some third-party compression tools have file size limitations, and cannot access files larger than 4 GB.

Chinese characters stored as ? in an Oracle database

When you configure an Oracle database to use the ZHS16GBK character set, some Chinese characters or symbols that are encoded with the GBK character set are incorrectly stored as a question mark (?). This problem is caused by an incompatibility of the GBK undefined code range Unicode mapping between the Oracle ZHS16GBK character set and the IBM GBK converter. To fix this problem, use a new code page, MS936A, by including the following system property when you start the JVM:

```
-Dfile.encoding=MS936A
```

For IBM WebSphere Application Server users, this problem might occur when web applications that use JDBC configure Oracle as the WebSphere Application Server data source. To fix this problem, use a new code page, MS936A, as follows:

1. Use the following system property when you start the JVM:

```
-Dfile.encoding=MS936A
```
2. Add the following lines to the `WAS_HOME/properties/converter.properties` file, where `WAS_HOME` is your WebSphere Application Server installation directory.

```
GBK=MS936A  
GB2312=MS936A
```

Issues with the XL TXE-J XSLT compiler

A low split limit might cause compilation errors.

Avoid calling Java extension functions that have side effects because the order of execution is not guaranteed.

Versions of Ant before 1.7.0 do not work with the XL TXE-J compiler. Instead, use the XSLT4J interpreter by running Java with the following system property:

- `-Djavax.xml.transform.TransformerFactory=org.apache.xalan.processor.TransformerFactoryImpl.`

Large page request fails

There is no error message issued when the JVM is unable to honor the `-Xlp` request.

There are a number of reasons why the JVM cannot honor a large page request. For example, there might be insufficient large pages available on the system at the time of the request. To check whether the `-Xlp` request was honored, you can review the output from `-verbose:gc`. Look for the attributes `requestedPageSize` and `pageSize` in the `-verbose:gc` log file. The attribute `requestedPageSize` contains the value specified by `-Xlp`. The attribute `pageSize` is the actual page size used by the JVM.

Unexpected CertificateException

IBM SDK for z/OS, V6 service refresh 13 fix pack 1 and later releases contain a security enhancement to correctly validate certificates on jar files of applications. After upgrading, a `CertificateException` occurs for any applications in one of the following scenarios:

- The application jar is not properly signed.
- The application jar has incorrect certificates.
- A certificate in the certificate chain is revoked.

To avoid these exceptions, make sure that your application jars are signed with valid certificates before you upgrade from an earlier release. This issue relates to APAR IV38456.

Unexpected application errors with RMI

If your application uses RMI and you experience unexpected errors after updating to IBM SDK for z/OS, V6 service refresh 13 fix pack 2, or later releases, the problem might be associated with a change to the default value of the RMI property `java.rmi.server.useCodebaseOnly`. For more information, see <http://docs.oracle.com/javase/7/docs/technotes/guides/rmi/enhancements-7.html>.

Unexpected XSLT error on extension elements or extension functions when Java security is enabled

From IBM SDK for z/OS, V6 service refresh 14, any attempt to use extension elements or extension functions when Java security is enabled, results in a `javax.xml.transform.TransformerException` error during XSLT processing. This change in behavior is introduced to enhance security.

The following XSLT message is generated when extension functions are used: Use of the extension function '<method name>' is not allowed when Java security is enabled. To override this, set the `com.ibm.xtq.processor.overrideSecureProcessing` property to true. This override only affects XSLT processing.

The following XSLT message is generated when extension elements are used: Use of the extension element '<element name>' is not allowed when Java security is enabled. To override this, set the `com.ibm.xtq.processor.overrideSecureProcessing` property to `true`. This override only affects XSLT processing.

To allow extensions when Java security is enabled, set the **`com.ibm.xtq.processor.overrideSecureProcessing`** system property to `true`. For more information about this system property, see “-Dcom.ibm.xtq.processor.overrideSecureProcessing” on page 77.

Support for virtualization software

The IBM SDK for Java is tested with a number of virtualized server products.

The following virtualization software is tested with the latest release of the IBM SDK for z/OS, V6:

Table 7. Virtualization software tested for the IBM SDK for Java

Vendor	Architecture	Server virtualization	Version
IBM	System z	PR/SM™	z10™, z11, z196, zEC12
IBM	System z	z/VM®	6.1, 6.2
IBM	POWER	PowerVM® Hypervisor	Power6, Power7
VMware	x86-64	VMware ESX and ESXi Server	4.1, 5.0
Red Hat	x86-64	Red Hat Enterprise Virtualization (RHEV)	2.1, 3.0
SUSE	x86-64	SUSE KVM	SLES 11
Microsoft	x86-64	Hyper-V	Server 2012

Notices

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

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

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

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1758
U.S.A.

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

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law:

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

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

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web

sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

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

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

JIMMAIL@uk.ibm.com

[Hursley Java Technology Center (JTC) contact]

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

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

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

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

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

COPYRIGHT LICENSE:

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

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows:

© (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. _enter the year or years_.

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

Privacy Policy Considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, see: (i) IBM’s Privacy Policy at <http://www.ibm.com/privacy> ; (ii) IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> (in particular, the section entitled “Cookies, Web Beacons and Other Technologies”); and (iii) the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

Trademarks

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

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.



Printed in USA