



BEARocket

Developing Java Applications

Version 5.0 Service Pack 2
June 2005

Copyright

Copyright © 2005 BEA Systems, Inc. All Rights Reserved.

Restricted Rights Legend

This software and documentation is subject to and made available only pursuant to the terms of the BEA Systems License Agreement and may be used or copied only in accordance with the terms of that agreement. It is against the law to copy the software except as specifically allowed in the agreement. This document may not, in whole or in part, be copied, photocopied, reproduced, translated, or reduced to any electronic medium or machine readable form without prior consent, in writing, from BEA Systems, Inc.

Use, duplication or disclosure by the U.S. Government is subject to restrictions set forth in the BEA Systems License Agreement and in subparagraph (c)(1) of the Commercial Computer Software-Restricted Rights Clause at FAR 52.227-19; subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013, subparagraph (d) of the Commercial Computer Software--Licensing clause at NASA FAR supplement 16-52.227-86; or their equivalent.

Information in this document is subject to change without notice and does not represent a commitment on the part of BEA Systems. THE SOFTWARE AND DOCUMENTATION ARE PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND INCLUDING WITHOUT LIMITATION, ANY WARRANTY OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. FURTHER, BEA Systems DOES NOT WARRANT, GUARANTEE, OR MAKE ANY REPRESENTATIONS REGARDING THE USE, OR THE RESULTS OF THE USE, OF THE SOFTWARE OR WRITTEN MATERIAL IN TERMS OF CORRECTNESS, ACCURACY, RELIABILITY, OR OTHERWISE.

Trademarks or Service Marks

BEA, BEA JRockit, BEA Liquid Data for WebLogic, BEA WebLogic Server, Built on BEA, Jolt, JoltBeans, SteelThread, Top End, Tuxedo, and WebLogic are registered trademarks of BEA Systems, Inc. BEA Builder, BEA Campaign Manager for WebLogic, BEA eLink, BEA Manager, BEA MessageQ, BEA WebLogic Commerce Server, BEA WebLogic Enterprise, BEA WebLogic Enterprise Platform, BEA WebLogic Enterprise Security, BEA WebLogic Express, BEA WebLogic Integration, BEA WebLogic Java Adapter for Mainframe, BEA WebLogic JDriver, BEA WebLogic JRockit, BEA WebLogic Log Central, BEA WebLogic Personalization Server, BEA WebLogic Platform, BEA WebLogic Portal, BEA WebLogic Server Process Edition, BEA WebLogic WorkGroup Edition, BEA WebLogic Workshop, and Liquid Computing are trademarks of BEA Systems, Inc. BEA Mission Critical Support is a service mark of BEA Systems, Inc. All other company and product names may be the subject of intellectual property rights reserved by third parties.

All other trademarks are the property of their respective companies.

Contents

Introduction

Recommended Coding Practices

Read the Relevant Specifications	2-4
Example 1: Reflection	2-4
Example 2: Reflection Revisited	2-4
Example 3: Serialization	2-5
Never Use Deprecated Unsafe Methods	2-5
Minimize the Use of Finalizers	2-5
Don't Depend on Thread Priorities	2-5
Don't Use Internal sun.* or COM.jrockit.* Classes	2-6
Override java.Object.hashCode for User Defined Classes When Using java.util.Hashtable	2-6
Do Careful Thread Synchronization	2-6
Expect Only Standard System Properties	2-6
Minimize the Number of Java Processes	2-7
Avoid Calling System.gc()	2-7

Troubleshooting

An Application Does Not Run	3-10
Slow-to-Start Applications	3-10
Process Counter Does Not Initialize	3-11
Large Memory Consumption	3-11

Slow Performance vis-a-vis HotSpot	3-12
Randomly Appearing Bugs	3-12
BEA JRockit JVM Throws Errors HotSpot Does Not Throw	3-12
Slow Performance in Development Mode	3-13
BEA JRockit JVM Does Not Run Jakarta Tomcat as a Windows Service	3-13
Other Frequently Asked Questions	3-14

Profiling and Debugging with BEA JRockit

Profiling BEA JRockit	4-15
Using JVMPI	4-15
How JVMPI Works	4-16
Changing the JVMPI Default Behavior	4-16
Additional JVMPI Documentation	4-18
Profiling with the HPROF Profiling Agent	4-18
HPROF Documentation	4-18
Debugging with BEA JRockit	4-18
Java Virtual Machine Debugger Interface (JVMDI)	4-19
How JVMDI Works	4-19
JVMDI Documentation	4-19

Migrating to BEA JRockit

About Application Migration	5-21
Why Migrate?	5-21
Migration Restrictions	5-22
Migration Support	5-22
Migration Procedures	5-22
Environment Changes	5-23
Other Tips	5-23

Tuning BEA JRockit JVM for Your Application	5-23
Testing the Application	5-24
Why Test?	5-24
How to Test.	5-24
Submitting Migration Tips	5-24

Introduction

Welcome to *Developing Java Applications*, a guide for Java developers creating Java applications and then migrating them from third-party JVMs to BEA JRockit JVM. This document contains the following information:

- The best coding practices developers should follow to ensure optimal performance with any JVM, particularly BEA JRockit 5.0 JDK. See [Recommended Coding Practices](#).
- A detailed troubleshooting guide that will lead you through the solutions to some of the more common problems developers have encountered with BEA JRockit JDK. See [Troubleshooting](#).
- A description of the debugging and profiling tools available with BEA JRockit SDK. See [Profiling and Debugging with BEA JRockit](#).
- Procedures for migrating applications developed on a third-party JVM to BEA JRockit JDK. See [Migrating to BEA JRockit](#).

Introduction

Recommended Coding Practices

This section contains guidelines for writing applications to run on BEA JRockit. This information provided here is in no way complete; it merely helps you avoid some common pitfalls. BEA Systems does not want to compromise Java's "write once run everywhere" notion. On the contrary, this sections highlights guidelines that are valid for any Java program. They are, however, especially important when switching between JVMs in general and between Sun Microsystem's HotSpot JVM and BEA JRockit JVM in particular.

The best coding practices are summarized in the following subjects:

- [Read the Relevant Specifications](#)
- [Never Use Deprecated Unsafe Methods](#)
- [Minimize the Use of Finalizers](#)
- [Don't Depend on Thread Priorities](#)
- [Don't Use Internal sun.* or COM.jrockit.* Classes](#)
- [Override java.Object.hashCode for User Defined Classes When Using java.util.Hashtable](#)
- [Do Careful Thread Synchronization](#)
- [Expect Only Standard System Properties](#)
- [Minimize the Number of Java Processes](#)
- [Avoid Calling System.gc\(\)](#)

Read the Relevant Specifications

Read the Java language specification and Java API specification carefully and do not rely on unspecified behavior.

A BEA JRockit JVM is based on a number of specifications; for example, The Java Virtual Machine Specification and the Java API Specification. You should be aware that many implementations of these specifications exist: BEA JRockit JVM is one. You should never expect any particular behavior that is *not* specified in one of these documents. Unspecified behavior might differ between the Sun JVM and BEA JRockit JVM. Note too that behavior is sometimes different between individual releases of the Sun JVM and can also change between releases of BEA JRockit JVM.

You can find these specifications at the following sites:

- The Java Language Specification

http://java.sun.com/docs/books/jls/second_edition/html/j.title.doc.html

- Java 2 Platform API Specification,

<http://java.sun.com/j2se/1.5.0/docs/api/index.html>

- The Java Virtual Machine Specification

<http://java.sun.com/docs/books/vmspec/2nd-edition/html/VMSpecTOC.doc.html>

The specifications are written to give JVM vendors freedom to optimize their JVMs, and therefore they leave certain behavior unspecified. You should understand, however, that numerous parts of the specifications mentioned above are unspecified. The following examples describe three of these unspecified elements.

Example 1: Reflection

The Java API Specification of the method `getMethods()` on the `java.lang.Class` class clearly states: “The elements in the array returned are not sorted and are not in any particular order.”

Example 2: Reflection Revisited

The `toString()` method of the `java.lang.reflect.Method` might include the access modifier `native`. Therefore, you should not rely on the result of this call to be equal between JVM implementations. Some classes in the Java API specification are implemented as `native` either by BEA JRockit JVM and the Sun JVM. There is no guarantee that a native implementation on one JVM has to be native on another one.

Example 3: Serialization

The Java API Specification of the method `defaultReadObject()` of the `java.lang.ObjectInputStream` class does not specify the order in which fields are de-serialized; hence no such order can be expected.

Never Use Deprecated Unsafe Methods

The deprecated methods

- `java.lang.Thread.stop`
- `java.lang.Thread.suspend`
- `java.lang.Thread.resume`
- `java.lang.Runtime.runFinalizersOnExit`

are inherently unsafe, and should never be used. For more information see:

<http://java.sun.com/j2se/1.5.0/docs/guide/misc/threadPrimitiveDeprecation.html>

Minimize the Use of Finalizers

Finalizers are often error prone since they often implicitly depend on the order of execution. This order differs amongst JVMs, and between consecutive runs on the same JVM. Using finalizers is also inherently bad for performance since it imposes an additional burden on the memory management system that needs to handle execution of finalizers and let objects live longer. For more information on using—and not using—finalizers, please refer to:

<http://access1.sun.com/techarticles/weak.references.html>

<http://www.memorymanagement.org/glossary/r.html#reference.object>

Don't Depend on Thread Priorities

Be careful when using `java.lang.Thread.setPriority`. Depending on thread priorities might lead on unwanted or unexpected results since the scheduling algorithm might choose to starve lower priority threads of CPU time and never execute them. Furthermore the result might differ between operating systems and JVMs.

The Java API specification states that “Every thread has a priority. Threads with higher priority are executed in preference to threads with lower priority.”

The priority set by the `setPriority()` method is a parameter that might be used in the thread-scheduling algorithm, which shares CPU execution time between executing threads. This

algorithm might be controlled either by the JVM or by the operating system. It is important to be aware of the fact that this algorithm normally differs between operating systems and that the algorithm might change between releases of both the operating system and the JVM. For BEA JRockit JVM native threads, the algorithm is implemented by the operating system.

Don't Use Internal `sun.*` or `COM.jrockit.*` Classes

The classes that BEA JRockit JDK includes fall into package groups `java.*`, `javax.*`, `org.*`, `sun.*` and `COM.jrockit.*`. All but the `sun.*` and `COM.jrockit` packages are a standard part of the Java platform and will be supported into the future. In general, non-standard packages, which are outside of the Java platform, can be different across JVM vendors and OS platforms (Windows, Linux, and so on) and can change at any time without notice with JDK versions. Programs that contain direct usage of the `sun.*` and `COM.jrockit.*` packages are not 100% Pure Java.

For more information, please refer to the note about `sun.*` packages at:

<http://java.sun.com/products/jdk/faq/faq-sun-packages.html>

Override `java.Object.hashCode` for User Defined Classes When Using `java.util.Hashtable`

On BEA JRockit JVM, the current default implementation of `hashCode` returns a value for the object determined by the JVM. The value is created using the memory address of the object. However, because this value can be reused if the object is moved during garbage collection, it is possible to obtain the same hash code for two different objects. Also, two objects that represent the same value are guaranteed to have the same hash code only if they are the exact same object. This implementation is not particularly useful for hashing; therefore, derived classes should override `hashCode()`.

Do Careful Thread Synchronization

Make sure that you synchronize threads that access shared data. Synchronization bugs often appear when changing JVMs because the implementation of locks, garbage collection, thread scheduling and so on, might differ significantly.

Expect Only Standard System Properties

When implementing `java.lang.System.getProperties()` or `java.lang.System.getProperty()`, you should only depend on standard system properties

being returned; different VMs may return a different set of extended properties. Non-standard properties should not be returned.

When the JVM starts, it inserts a number of standard properties into the system properties list. These properties, and the meaning of their values, are listed in the Java API specification. Do not expect any other non-standard properties.

<http://java.sun.com/j2se/1.5.0/docs/api/java/lang/System.html>

Minimize the Number of Java Processes

When designing applications there is sometimes a choice between running several processes, i.e. JVM instances versus running several threads or thread groups within a single process, i.e. JVM instance. If possible, it is more effective to use as few JVM instances as possible per physical machine.

Avoid Calling `System.gc()`

Don't call `java.lang.System.gc()`. This method behaves differently with BEA JRockit JVM than with other JVMs. Instead of doing a complete garbage collection, as with the Sun JVM and others, when called by an application running with BEA JRockit JVM, `System.gc()` behaves depending upon the garbage collector already in use:

- If you're using a generational copy collector, `System.gc()` does a collection in their nursery.
- For all other collectors, it does a collection only if one is needed. In other words, it does nothing special on the call.

Note: If you must call `System.gc()`, you can override BEA JRockit JVM's behavior by using the command line option `-XXfullsystemgc`.

The BEA JRockit JVM garbage collector will generally do a much better job of deciding when to do garbage collection than will `System.gc()`. If you are having problems with memory usage, pause times for garbage collection, and so on, you are better off configuring the BEA JRockit JVM memory management system appropriately. See [Tuning BEA JRockit JVM](#).

Troubleshooting

While the process of switching to BEA JRockit 1.4.2 JVM from another JVM is relatively easy and generally problem-free, you might encounter some known issues while or after making this switch. This section describes some of those issues and describes some simple workarounds. The issues that might occur are:

- [An Application Does Not Run](#)
- [Slow-to-Start Applications](#)
- [Process Counter Does Not Initialize](#)
- [Large Memory Consumption](#)
- [Slow Performance vis-a-vis HotSpot](#)
- [Randomly Appearing Bugs](#)
- [BEA JRockit JVM Throws Errors HotSpot Does Not Throw](#)
- [Slow Performance in Development Mode](#)
- [BEA JRockit JVM Does Not Run Jakarta Tomcat as a Windows Service](#)
- [Other Frequently Asked Questions](#)

An Application Does Not Run

I cannot get my favorite Java application to run on BEA JRockit JVM. What am I doing wrong?

Many problems with running applications on BEA JRockit JVM is due to erroneous environment variables or non-standard startup options.

Start by ensuring that your environment variables are set up correctly. Make sure that you have set your `JAVA_HOME` environment variable correctly, i.e. set to the directory where BEA JRockit JVM has been installed, and that "`%JAVA_HOME%\bin`" is available in your `PATH` environment variable before any other directory where any version of `java.exe` may exist. When running applications as Windows services it is crucial that you set these environment variables system wide. To do this:

1. Open the Start menu and select Settings>Control Panel>System
2. Select the Advanced tab.
3. Click Environment Variables.
4. To set system wide environment variables you must edit in the lower part of this dialog box, labeled System variables.

Applications are often started via scripts. Make sure that none of the startup scripts includes non-standard startup options for java. See [Tuning BEA JRockit JVM](#) for complete documentation of standard and non-standard options.

Slow-to-Start Applications

Why does it take longer for my applications to start with BEA JRockit JVM?

Java programs are compiled into byte code by a Java compiler. Many JVMs, including the Sun JVM, interprets this byte code each time it is executed. BEA JRockit JVM, however, uses code generation technology to generate native machine code from the byte code. This is sometimes called Just-In-Time (JIT) compilation. The code generation step imposes an initial time penalty before execution. Normally, the subsequent execution of the code is faster than interpreting the byte code. BEA JRockit JVM is optimized for server applications that normally run for long periods of times. The initial time penalty is normally negligible in comparison to the performance gains of code generation over time.

Process Counter Does Not Initialize

Sometimes, when running BEA JRockit, I encounter a `NotAvailableException` when the console tries to connect to JRockit or if the program itself tries to access the CPU load counters.

Occasionally, the process counter does not initialize. This happens only on Windows installations where either the security settings are such that the Performance Data Helper (PDH; a Windows API that reads performance metrics from the operating system) process counter can't be read or where, for reasons unknown, the PDH process counter is simply turned off. This will deny you the rights to look at the process counter and throw the error. You will receive this message:

```
JRockit] WARNING: Could not initialize the virtualbytes counter, some
functionality in jrockit.management.Memory will not be available. Message
was: failed to create counter query. String was: (null)\Virtual Bytes

[JRockit] WARNING: Could not initialize the JVM process load counter, CPU
load generated by the JVM will not be available. Message was: failed to create
counter query. String was: (null)\% Processor Time
```

Why the process counter is turned off is unknown; however, should you encounter this situation, you can turn it on again by following the instructions at this location:

<http://www.microsoft.com/windows2000/techinfo/reskit/tools/existing/exctrlst-o.asp>

If that doesn't work, check your security settings and then make sure the performance counters can be read using the Windows `perfmon` tool, running as the same user that the JVM process is running as.

Large Memory Consumption

Why does my applications consume more memory when running on BEA JRockit JVM?

The Java programming language relies on a mechanism called garbage collection (GC) to free memory when it is no longer being used. There is no equivalence to the `delete` operator in the C++ programming language or the `free` function in the C programming language. Any Java virtual machine must include a garbage collector that handles the task of finding unreferenced objects, possibly invoke their finalizers and free the memory used to hold their state.

The BEA JRockit JVM garbage collectors are described in [Selecting and Running a Memory Management System](#) in *Using BEA JRockit DK*. Generally, the BEA JRockit JVM garbage

collection implementations trade high memory usage for speed and minimal program wide halts; that is, acquiring system wide locks.

Slow Performance vis-a-vis HotSpot

I have a script/program that use BEA JRockit JVM for certain tasks. Why is it slower than when I use the Sun JVM, HotSpot?

This might be related to the [Slow-to-Start Applications](#) above. Scripts or other programs may start many Java processes and may therefore experience bad performance compared to the Sun JVM, since BEA JRockit JVM has a code generation penalty when starting up. When starting many Java processes and running them only for a short time, this penalty can become significant.

Randomly Appearing Bugs

Why does my application have randomly appearing bugs when running on BEA JRockit JVM that it doesn't have when running on the Sun JVM?

You may be experiencing synchronization bugs in your application. It is not uncommon that such bugs are revealed when switching JVMs. The JVM specification and the Java language specification leaves plenty of room for optimization that may cause unsynchronized access to shared data, to cause different behavior on different JVMs.

For more information see:

http://java.sun.com/docs/books/jls/second_edition/html/memory.doc.html

<http://java.sun.com/docs/books/vmspec/2nd-edition/html/Threads.doc.html>

BEA JRockit JVM Throws Errors HotSpot Does Not Throw

Why is BEA JRockit JVM throwing `IllegalAccessError`, `ClassFormatError`, `IncompatibleClassChangeError` or other `LinkageError` exceptions when the Sun JVM is not?

Verification ensures that the binary representation of a class or interface is structurally correct. For example, it checks that every instruction has a valid operation code; that every branch instruction branches to the start of some other instruction, rather than into the middle of an instruction; that every method is provided with a structurally correct signature; and that every instruction obeys the type discipline of the Java programming language.

If an error occurs during verification, then an instance of the following subclass of class `LinkageError` will be thrown at the point in the program that caused the class to be verified

Example: Using JTidy throws `IllegalAccessError`

In an early version of JTidy from Apache Software Foundation, the compiler had incorrectly inlined a reference to a private variable belonging to an outer class. This caused an exception to be thrown since BEA JRockit JVM does stricter verification than the Sun JVM is. The old `Tidy.jar` should be replaced with the new and correctly compiled version.

Slow Performance in Development Mode

Why is BEA JRockit JVM slower when BEA WebLogic Server is running in development mode?

When WebLogic Server is started in development mode, BEA JRockit JVM is by default started with the `-xdebug` option. This makes BEA JRockit JVM run with some overhead.

Note: This option is purely for diagnostics use and should therefore not be used in a production type environment.

BEA JRockit JVM Does Not Run Jakarta Tomcat as a Windows Service

I cannot get BEA JRockit JVM to run Jakarta Tomcat as a Windows service. What am I doing wrong?

The quick answer to this problem is: if you are using `jk_nt_service`, do everything that you need to do for the Sun JVM, then exchange the non-standard Sun JVM `-xrs` startup option with the non-standard BEA JRockit JVM `-xnohup` in the `wrapper.properties` configuration file. The rest of this answer is a slightly more detailed description of this.

1. First make sure that you have completed all the tasks concerning the environment variables in “I cannot get my favorite Java application to run on BEA JRockit JVM. What am I doing wrong?” above.

Many people use the `jk_nt_service` windows service wrapper to run java applications; for example, Jakarta Tomcat, as a Windows service (see <http://members.ozemail.com.au/~lampante/howto/tomcat/iisnt/>). Independently of what Windows service you may be using you must make sure that it is not using any non-standard startup options. When using `jk_nt_service`, the startup is defined in:

```
<tomcat install dir>\conf\jk\wrapper.properties
```

2. Make sure that you set the three properties `wrapper.tomcat_home`, `wrapper.java_home` and `wrapper.cmd_line` are set accordingly:

- `wrapper.tomcat_home` must be set with the installation directory of tomcat
- `wrapper.java_home` must be set to the same value as the `JAVA_HOME` environment variable.
- The property `wrapper.cmd_line` defines the startup command. At the time of writing, this property should be set to:

```
wrapper.cmd_line=$(wrapper.javabin) -Xnohup
-Djava.security.policy=="$(wrapper.tomcat_policy)"
-Dtomcat.home="$(wrapper.tomcat_home)" -classpath
$(wrapper.class_path) $(wrapper.startup_
```

for BEA JRockit JVM. Normally this command includes a non-standard option to stop the JVM from shutting down the process when a user logs off. For BEA JRockit JVM this non-standard option is `-Xnohup`, for the Sun JVM it is `-Xrs`.

Other Frequently Asked Questions

What is a Java virtual machine?

As described by the [Java Virtual Machine Specification](#):

“The Java virtual machine is the cornerstone of the Java and Java 2 platforms. It is the component of the technology responsible for its hardware- and operating system- independence, the small size of its compiled code, and its ability to protect users from malicious programs.”

What is the difference between the Sun JVM and BEA JRockit JVM?

The most well know JVM is the implementation from Sun. The Sun JVM is called HotSpot. The Sun JVM is shipped in the Java Developer's Kit (JDK) and Java Runtime Environment (JRE) from Sun.

The BEA JRockit JVM from BEA systems is optimized for reliability and performance for server side applications. To achieve this, BEA JRockit JVM uses technologies such as code generation, hot spot detection, code optimization, advanced garbage collection algorithms and tight operating system integration.

Should I write my applications differently for BEA JRockit JVM?

No! You should not write your applications in any other way for BEA JRockit JVM than you should for any other JVM. You should, however, design and implement your applications well in order for them to run well on BEA JRockit JVM.

Profiling and Debugging with BEA JRockit

BEA JRockit SDK includes the JVM profiling interface (JVMPPI) and JVM debugging interface (JVMDI) which enable Java applications to interact with the JVM to assist with profiling and debugging activities. While developers will need to implement these interfaces within their application code, users' exposure to JVMPPI and JVMDI will usually be through the profiling and debugging tools they select for the applications they are running.

This section includes information on the following subjects:

- [Profiling BEA JRockit](#)
- [Profiling with the HPROF Profiling Agent](#)

Profiling BEA JRockit

You can use any number of third-party profiling tools to profile BEA JRockit performance. This section describes how to use the Java Virtual Machine Profiler Interface (JVMPPI) to facilitate using those tools.

Using JVMPPI

The JVM Profiler Interface allows you to use third-party profiling tools with BEA JRockit SDK.

Warning: This interface is an experimental feature in the Java 2 JDK and is not yet a standard profiling interface.

How JVMPI Works

JVMPI is a two-way function call interface between the Java virtual machine and an in-process profiler agent. On one hand, the VM notifies the profiler agent of various events, corresponding to, for example, heap allocation, thread start, and so on. Concurrently, the profiler agent issues controls and requests for more information through the JVMPI. For example, the profiler agent can turn on/off a specific event notification, based on the needs of the profiler front-end.

The profiler front-end may or may not run in the same process as the profiler agent. It may reside in a different process on the same machine, or on a remote machine connected via the network. The JVMPI does not specify a standard wire protocol. Tools vendors may design wire protocols suitable for the needs of different profiler front-ends.

A profiling tool based on JVMPI can obtain a variety of information such as heavy memory allocation sites, CPU usage hot-spots, unnecessary object retention, and monitor contention, for a comprehensive performance analysis.

JVMPI supports partial profiling; that is, a user can selectively profile an application for certain subsets of the time the VM is up and can also choose to obtain only certain types of profiling information.

Note: JVMPI supports only one agent per VM.

Changing the JVMPI Default Behavior

Use the following option to modify the JVMPI default behavior:

```
-Xjvmpi [[:<argument1>=<value1>[ ,<argumentN>=<valueN>]]
```

When BEA JRockit runs with a profiling agent attached, by default a number of events are enabled that can create significant overhead. Since JVMPI doesn't require all of these events to be sent,

you can disable them by setting the `-xjvmpi` flag. Use the arguments listed in [Table 4-1](#) to modify the default behavior.:

Table 4-1 Command Line Arguments for `-xjvmpi`

Argument	Description
<code>entryexit=off on (default on)</code>	Setting this to <code>off</code> disables the following method entry and exit events sent by JVMPI: <ul style="list-style-type: none">• <code>JVMPI_EVENT_METHOD_ENTRY</code>• <code>JVMPI_EVENT_METHOD_ENTRY2</code>• <code>JVMPI_EVENT_METHOD_EXIT</code>
<code>allocs=off on (default on)</code>	Setting this to <code>off</code> disables these object allocation and free events: <ul style="list-style-type: none">• <code>JVMPI_EVENT_OBJECT_ALLOC</code>• <code>JVMPI_EVENT_OBJECT_MOVE</code>• <code>JVMPI_EVENT_OBJECT_FREE</code>• <code>JVMPI_EVENT_ARENA_NEW</code>• <code>JVMPI_EVENT_ARENA_DELETE</code>
<code>monitors=off on (default on)</code>	Setting this to <code>off</code> disables these monitor contention events: <ul style="list-style-type: none">• <code>JVMPI_EVENT_RAW_MONITOR_CONTENTENDED_ENTER</code>• <code>JVMPI_EVENT_RAW_MONITOR_CONTENTENDED_ENTERED</code>• <code>JVMPI_EVENT_RAW_MONITOR_CONTENTENDED_EXIT</code>• <code>JVMPI_EVENT_MONITOR_CONTENTENDED_ENTER</code>• <code>JVMPI_EVENT_MONITOR_CONTENTENDED_ENTERED</code>• <code>JVMPI_EVENT_MONITOR_CONTENTENDED_EXIT</code>• <code>JVMPI_EVENT_MONITOR_WAIT</code>• <code>JVMPI_EVENT_MONITOR_WAITED</code>
<code>arenadelete=off on (default off)</code>	Setting this to <code>on</code> will enable the <code>JVMPI_EVENT_ARENA_DELETE</code> event. This event is suppressed by default to be compatible with Sun's VM which does not send this event. The event can be enabled if a profiler wishes to receive the event.

Additional JVMPI Documentation

As JVMPI is an experimental interface, Sun Microsystems provides the documentation for tools vendors who have an immediate need for profiling hooks in the Java VM. You can find this documentation at:

<http://java.sun.com/j2se/1.5.0/docs/guide/jvmpi/index.html>

Profiling with the HPROF Profiling Agent

An ancillary component to JVMPI is the HPROF profiling agent, which is shipped with the Java 2 JDK. HPROF is a dynamically-linked library that interacts with the JVMPI and writes profiling information either to a file or to a socket. You can then process that information by using a profiler front-end tool.

HPROF displays such information as CPU usage, heap allocation statistics, and monitor contention profiles. It can also report complete heap dumps and states of all the monitors and threads in BEA JRockit.

To run HPROF, use the `-Xrunhprof` command at startup; for example:

```
java -Xrunhprof ClassToProfile
```

Depending on the type of profiling requested, HPROF instructs the virtual machine to send it the relevant JVMPI events and processes the event data into profiling information. For example, the following command obtains the heap allocation profile:

```
java -Xrunhprof:heap=sites ToBeProfiledClass
```

HPROF Documentation

HPROF is distributed as part of the J2SE JDK. You can find complete documentation for this feature at:

<http://java.sun.com/j2se/1.5.0/docs/guide/jvmpi/jvmpi.html#hprof>

Debugging with BEA JRockit

This section describes the interface by which debugging tools can interface with BEA JRockit to debug Java applications.

Java Virtual Machine Debugger Interface (JVMDI)

JVMDI is a low-level debugging interface used by debuggers and other programming tools. It allows you to inspect the state and to control the execution of applications running in the BEA JRockit JVM.

JVMDI describes the functionality a JVM provides to enable debugging of Java applications running within the JVM. JVMDI defines the services a JVM must provide for debugging. JVMDI services include requests for information (for example, current stack frame), actions (set a breakpoint), and notification (when a breakpoint has been hit).

How JVMDI Works

JVMDI is a two-way interface:

- The JVMDI client can be notified of interesting occurrences through events.
- The JVMDI can query and control the application through many different functions, either in response to events or independent of them.

JVMDI clients run in the same VM as the application being debugged and access JVMDI through a native interface. The native, in-process interface allows maximum control with minimal intrusion of a debugging tool. Typically, JVMDI clients are relatively compact. They can be controlled by a separate process that implements the bulk of a debugger's functionality without interfering with the target application's normal execution.

JVMDI Documentation

Sun Microsystems provides complete reference documentation for the Java Platform Debug Architecture and JVMDI. For more information, go to:

<http://java.sun.com/j2se/1.4.1/docs/guide/jpda/index.html>

Migrating to BEA JRockit

This section describes how to migrate Java applications developed on another JVM to BEA JRockit so that they perform to their optimal capability when running on BEA WebLogic Server. It contains information on the following subjects:

- [About Application Migration](#)
- [Migration Procedures](#)
- [Testing the Application](#)
- [Submitting Migration Tips](#)

About Application Migration

Migrating an application to BEA JRockit JVM is a relatively simple process, requiring some minor environmental changes and following some simple coding guidelines. This section provides instructions and tips to successfully completing this simple process. It also describes some of the benefits and possible problems you might encounter during migration and it discusses some best J2SE coding practices for you to follow to ensure that your application runs successfully once it is running on BEA JRockit.

Why Migrate?

BEA JRockit JVM is the default JVM shipped with BEA WebLogic Server. Although there are other JVMs available on the market today that you can use to develop Java applications, BEA

Systems recommends that you use BEA JRockit JVM as the production JVM for any application deployed on WebLogic Server.

Migration Restrictions

Migration is available only for Intel-based Windows systems and Linux systems. For a list of supported platforms, please refer to:

<http://edocs.bea.com/wljrookit/docs142/certif.html>

Migration Support

Should you experience any problems or find any bugs while attempting to migrate an application to BEA JRockit 8.1, please send us an e-mail at support@bea.com. We would appreciate if you could provide as much information as possible about the problem, for example:

- Hardware
- Operating system and its version
- The program you are attempting to migrate
- Stack dumps (if any)
- A small code example that will reproduce the error
- Copies of any *.dump and *.mdmp files (*.mdmp files are available only on Windows)

Migration Procedures

This section describes basic environmental and implementation changes necessary to migrate to BEA JRockit JVM from Sun Microsystems HotSpot JVM or any other third-party JVM. It includes information on the following subjects:

- [Environment Changes](#)
- [Other Tips](#)
- [Tuning BEA JRockit JVM for Your Application](#)
- [Testing the Application](#)
- [Submitting Migration Tips](#)

Environment Changes

To migrate from HotSpot (or any third-party JVM) to BEA JRockit JVM, you need to make the following changes to the files.

- Set the `JAVA_HOME` environmental variable in `<WEBLOGIC_HOME>/common/commEnv.cmd` (or `.sh`) to the appropriate path.
- Set the `JAVA_VENDOR` environmental variable in `<WEBLOGIC_HOME>/common/commEnv.cmd` (or `.sh`) to `BEA`.
- If you are using a start-up script, remove any Sun-specific (or other JVM provider) options from the start command line (like `-hotspot`). If possible, replace them with BEA JRockit-specific options; for example, `-jrockit`. Other flags that might need to be changed include `MEM_ARGS` and `JAVA_VM`.
- Change `config.xml` to point the default compiler setting(s) to the BEA JRockit `javac` compiler.

Other Tips

For information on other coding practices that will ensure a successful migration of your application to BEA JRockit JVM, please refer to [Recommended Coding Practices](#).

Tuning BEA JRockit JVM for Your Application

Once you've migrated your application to BEA JRockit JVM, you might want to tune the JVM for optimal performance. For example, you might want to specify a different start-up heap size or set custom garbage collection parameters. For more information on tuning BEA JRockit JVM, please refer to the [Tuning BEA JRockit JVM](#).

The non-standard options, that is, options preceded with `-x`, are critical tools for tuning a JVM at startup. These options change the behavior of BEA JRockit JVM to better suit the needs of different Java applications.

While all JVMs use non-standard options, the option names might not be the same from JVM to JVM; for example, while BEA JRockit JVM will accept the non-standard option `-xns` to set the nursery in generational concurrent and generational copying garbage collectors, Sun's HotSpot JVM uses the option `-XX:NewSize` to set this value.

If you are migrating an application to BEA JRockit, we recommend that you become familiar with the non-standard options available to you. For more information, please refer to [Command Line Options by Name](#).

You should also be aware that, being non-standard, non-standard options are subject to change at any time.

Testing the Application

Always test your application on BEA JRockit JVM before putting it into production. If you develop your application on the Sun JVM (HotSpot), you *must* test your application on JVM before you put it into production.

Why Test?

Some important reasons for testing are:

- Sometimes you might find bugs in your own program that don't occur on the Sun JVM; for example, synchronization problems.
- You might have used third party class libraries that are not 100% Java and rely on Sun-specific classes or behavior.
- You might have used third-party class files that are not correct. BEA JRockit has been known to enforce verification more rigorously than the Sun JVM.

How to Test

To test your application on BEA JRockit:

1. Run your application against any test scripts or benchmarks that are appropriate for that application.
2. If any problems occur, handle them as you normally would for the specific application.

Submitting Migration Tips

The migration tips discussed in this section represent an evolving list. Often, a successful migration to BEA JRockit depends as much upon the application being migrated as it does to the VMs being used. BEA Systems welcomes suggestions based upon your experiences with migrating applications to BEA JRockit. Feel free to submit any migration ideas or comments to the BEA JRockit SDK migration newsgroup at:

`jrockit.developer.interest.migration`