

KVM Porting Guide

KVM 1.0.3



Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303 USA
650 960-1300 fax 650 969-9131

Version 1.0.3
September 14, 2001

Copyright © 1998-2001 Sun Microsystems, Inc.

901 San Antonio Road, Palo Alto, CA 94303 USA

All rights reserved. Copyright in this document is owned by Sun Microsystems, Inc.

Sun Microsystems, Inc. (SUN) hereby grants to you at no charge a nonexclusive, nontransferable, worldwide, limited license (without the right to sublicense) under SUN's intellectual property rights that are essential to practice the K Virtual Machine (KVM) or J2ME CLDC Reference Implementation technology to use this document for internal evaluation purposes only. Other than this limited license, you acquire no right, title, or interest in or to the document and you shall have no right to use the document for productive or commercial use.

RESTRICTED RIGHTS LEGEND

Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-1(a).

SUN MAKES NO REPRESENTATIONS OR WARRANTIES ABOUT THE SUITABILITY OF THE SOFTWARE, EITHER EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT. SUN SHALL NOT BE LIABLE FOR ANY DAMAGES SUFFERED BY LICENSEE AS A RESULT OF USING, MODIFYING OR DISTRIBUTING THIS SOFTWARE OR ITS DERIVATIVES.

TRADEMARKS

Sun, Sun Microsystems, the Sun logo, Java, the Java Coffee Cup logo, JDK, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. UNIX® is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.

Contents

1. About This Document	1
1.1 Who Should Use This Document	1
1.2 Related Documentation	1
2. Introduction to KVM	3
2.1 K Virtual Machine (KVM)	3
2.2 Differences between KVM 1.0.3 and KVM 1.0.2	4
3. Compiler Requirements	5
4. Directory Structure	7
4.1 Overview	7
4.2 Directory <code>kvm/VmCommon</code>	8
4.3 Directory <code>kvm/VmExtra</code>	11
5. Required Port-Specific Files and Functions	13
5.1 File <code>machine_md.h</code>	13
5.2 File <code>main.c</code>	14
5.3 Runtime functions that require porting efforts	14
5.4 Required C library functions	15
6. Compilation Flags, Definitions and Macros	17

6.1	General compilation options	17
6.2	General system configuration options	18
6.3	Palm-specific system configuration options	19
6.4	Memory allocation settings	20
6.5	Garbage collection options	21
6.6	Class loading options	21
6.7	Interpreter execution options (KVM 1.0)	21
6.8	Interpreter execution options (KVM 1.0.2 and later)	22
6.8.1	Copying the virtual machine registers to local variables	23
6.8.2	Splitting uncommon bytecodes into a separate subroutine	24
6.8.3	Moving the test for thread rescheduling to branchpoints	25
6.8.4	Padding out the bytecode space	25
6.9	Java-level debugging options	25
6.10	VM-level debugging and tracing options	26
6.10.1	Including and excluding debugging code	26
6.10.2	Tracing options	27
6.11	Error handling macros	28
6.12	Miscellaneous macros and options	28
6.13	Overriding the compilation flags and other options from makefiles	29
7.	Virtual Machine Startup	31
7.1	Command line startup	31
7.2	Alternative VM startup strategies	33
7.3	Using a JAM (Java Application Manager)	33
8.	Class Loading, JAR Files, and Inflation	35
8.1	Generic class file loading	35
8.2	JAR file reader	37
8.2.1	Opening a JAR file	37
8.2.2	Closing a JAR file	37
8.2.3	Reading a JAR file entry	38

8.2.4	Reading multiple JAR file directory	38
8.3	Inflation	39
9.	64-bit Support	41
9.1	Setup	41
9.2	Alignment issues	43
10.	Native Code	45
10.1	Native code lookup tables	45
10.2	Implementing native methods	46
10.2.1	Include files	46
10.2.2	Accessing arguments from native methods	46
10.2.3	Returning a result from a native function	47
10.2.4	Shortcuts	47
10.2.5	Callbacks	48
10.2.6	Exception handling in native code	48
10.2.7	Useful functions in native code	48
10.2.8	Garbage collection issues	49
10.2.9	Initialization and reinitialization of global variables	54
10.3	Asynchronous native methods	54
10.3.1	Design of asynchronous methods	55
10.3.2	Implementation of asynchronous methods	57
11.	Event Handling	59
11.1	High-level description	59
11.1.1	Synchronous notification (blocking)	59
11.1.2	Polling in Java code	60
11.1.3	Polling in the bytecode interpreter	60
11.1.4	Asynchronous notification	61
11.2	Parameter passing and garbage collection issues	63
11.3	Implementation in KVM	63

11.4	Battery power conservation	65
12.	Class File Verification	67
12.1	Overview	67
12.2	Using the preverifier	69
12.2.1	General form	69
12.2.2	Preverifier options	69
12.2.3	Supported input file formats	70
12.2.4	JAR support in preverifier (since KVM 1.0.2)	71
12.3	Porting the verifier	72
12.3.1	Compiling the preverifier	72
13.	JavaCodeCompact (JCC)	73
13.1	JavaCodeCompact options	73
13.2	Porting JavaCodeCompact	74
13.3	Compiling JavaCodeCompact	75
13.4	JavaCodeCompact files	75
13.5	Executing JavaCodeCompact	76
13.6	Limitations	78
14.	Java Application Manager (JAM)	79
14.1	Using the JAM to install applications	80
14.1.1	Application launching	81
14.1.2	Application updating	82
14.2	JAM components	83
14.2.1	Security requirements	83
14.2.2	JAR file	83
14.2.3	Application Descriptor File	83
14.2.4	Network communication	85
14.3	Application lifecycle management	85
14.3.1	Termination of the KVM Task	85

14.4	Error handling	86
14.4.1	Error conditions	86
15.	Java-Level Debugging Support (KDWP)	89
15.1	Overall architecture	90
15.2	Debug Agent	91
15.2.1	Connections between a debugger and the KVM	91
15.2.2	Packet processing	93
15.3	Debugger support within KVM	93
15.3.1	Events	94
15.3.2	Breakpoints	95
15.3.3	Single stepping	96
15.3.4	Suspend and nosuspend options	96
15.4	Using the Debug Agent and the JPDA Debugger	97
15.4.1	Starting a debug session	97
15.4.2	Debugging example	99

Figures

- FIGURE 6-1 Error handling 28
- FIGURE 10-1 A native method 46
- FIGURE 10-2 Forbidding garbage collection 50
- FIGURE 10-3 Creating a handle 51
- FIGURE 10-4 Temporary roots 52
- FIGURE 10-5 Creating a global root 53
- FIGURE 10-6 Asynchronous implementation of `ReadBytes` 57
- FIGURE 10-7 Alternative asynchronous implementation of `ReadBytes` 58
- FIGURE 12-1 Two-phase verification 68
- FIGURE 15-1 Java-level debugging interface architecture 90
- FIGURE 15-2 Debugger and KVM connections 92

Tables

TABLE 3-1	Basic types	5
TABLE 3-2	Floating point types	6
TABLE 4-1	Distribution directories	7
TABLE 4-2	Files in VmCommon	8
TABLE 4-3	Files in VmExtra	11
TABLE 9-1	64-bit types	41
TABLE 9-2	Implementing longs	42
TABLE 9-3	Implementing both longs and floats	42
TABLE 10-1	Macros for popping arguments from the stack	47
TABLE 10-2	Macros for pushing arguments onto the stack	47
TABLE 10-3	Macros used in asynchronous methods	56

About This Document

This document provides information for porting the K Virtual Machine (KVM), version 1.0.3, to a new platform. KVM is a Java Virtual Machine implementation that is commonly used as the execution engine for J2ME CLDC (Java™ 2 Micro Edition, Connected Limited Device Configuration.)

1.1 Who Should Use This Document

This document is intended primarily to those individuals and companies who want to port Sun's reference implementation of the K Virtual Machine to a new platform. The document is useful also to those persons who want to learn more about the internal details of the KVM.

1.2 Related Documentation

The Java™ Language Specification (Java Series), Second Edition by James Gosling, Bill Joy, Guy Steele and Gilad Bracha. Addison-Wesley, 2000, ISBN 0-201-31008-2

The Java™ Virtual Machine Specification (Java Series), Second Edition by Tim Lindholm and Frank Yellin. Addison-Wesley, 1999, ISBN 0-201-43294-3

Programming Wireless Devices with the Java™ 2 Platform, Micro Edition (Java Series) by Roger Riggs, Antero Taivalsaari, and Mark VandenBrink. Addison-Wesley, 2001, ISBN 0-201-74627-1

Connected, Limited Device Configuration Specification, version 1.0, Java Community Process, Sun Microsystems, Inc.

http://jcp.org/aboutJava/communityprocess/jsr/jsr_030_j2melc.html

Mobile Information Device Profile Specification, version 1.0, Java Community Process, Sun Microsystems, Inc.

http://jcp.org/aboutJava/communityprocess/jsr/jsr_037_mid.html

Java 2 Platform Micro Edition (J2ME™) Technology for Creating Mobile Devices, A White Paper, Sun Microsystems, Inc.

<http://java.sun.com/products/cldc/wp/KVMwp.pdf>

KVM Debug Wire Protocol (KDWP) Specification, Sun Microsystems, Inc.

Introduction to KVM

2.1 K Virtual Machine (KVM)

KVM (also known as the *K Virtual Machine* or the *KJava Virtual Machine*) is a compact, portable Java™ virtual machine that has been designed specifically for small, resource-constrained devices such as cellular phones, pagers, personal organizers, mobile Internet devices, point-of-sale terminals, home appliances, and so forth.

The high-level design goal for the KVM team was to create the smallest possible “complete” Java virtual machine that would maintain all the central aspects of the Java programming language, and that would nevertheless run in a resource-constrained device with only a few tens or hundreds of kilobytes of available memory (hence the name K, for kilobytes). More specifically, KVM is designed to be

- small, with a static memory footprint of the virtual machine core starting from about 60 kilobytes (depending on the target platform and compilation options),
- clean and highly portable,
- modular and customizable,
- as “complete” and “fast” as possible without sacrificing the other design goals.

KVM is implemented in the C programming language, so it can easily be ported onto various platforms for which an ANSI C compiler is available. The virtual machine has been built around a straightforward bytecode interpreter with various compile-time flags and options for helping porting efforts and space optimization.

KVM has been developed as part of a larger effort to provide a modular, scalable architecture for the development and deployment of portable, dynamically downloadable and secure applications in consumer and embedded devices. This larger effort is called the *Java 2 Micro Edition* (also known as Java 2 ME or J2ME).

The K Virtual Machine is typically used as the implementation-level foundation for the following J2ME technology standards: *Connected, Limited Device Configuration* (CLDC) and *Mobile Information Device Profile* (MIDP). Further information on KVM, CLDC, MIDP and Java 2 Micro Edition in general is available in separate documents listed in Section 1.2, “Related Documentation.”

KVM is derived from a research system called *Spotless* developed originally at Sun Microsystems Laboratories. More information on the Spotless system is available in the Sun Labs technical report *The Spotless system: implementing a Java system for the Palm connected organizer*.

2.2 Differences between KVM 1.0.3 and KVM 1.0.2

KVM 1.0.3 is primarily a maintenance release that contains various bug fixes as well as some performance enhancements.

The main features of KVM 1.0.3 compared to KVM 1.0.2 include:

- Performance optimizations (redesigned monitor/synchronization operations, optimizations in code generated by the JavaCodeCompact tool, string optimizations)
- Enhancements to the event system, asynchronous I/O capabilities and the networking libraries to align the CLDC implementation better with the MIDP reference implementation
- Enhancements to the Java-level debugging interface, the preverifier and the compacting garbage collector
- More efficient JAM (Java Application Manager) implementation
- Command-line adjustable heap size for the Windows/Unix versions of the KVM

For most up-to-date information, refer to the release notes and KVM product website (<http://java.sun.com/products/kvm>).

Compiler Requirements

In order to be able to compile the KVM codebase, you must have a C compiler capable of compiling ANSI-compliant C files. Your compiler must define the basic C types as shown below in Table 3-1.

TABLE 3-1 Basic types

Type	Description
char	An 8-bit quantity. It can be signed or unsigned.
signed char	A signed 8-bit quantity.
unsigned char	An unsigned 8-bit quantity.
short	A signed 16-bit quantity.
unsigned short	An unsigned 16-bit quantity.
int	A signed quantity. It is either 16 or 32 bits.
unsigned int	A unsigned quantity. It is either 16 or 32 bits.
long	A signed 32-bit quantity.
unsigned long	An unsigned 32-bit quantity.
void *	A 32-bit pointer.

If your J2ME configuration or profile supports floating point numbers, your compiler must support the floating point types shown below in Table 3-2.

TABLE 3-2 Floating point types

Type	Description
float	A 32-bit floating point value.
double	A 64-bit floating point value.

All KVM implementations support the Java type `long`.¹ It is preferable that your compiler support 64-bit integers; however this is not a requirement. Porting the Java type `long` is discussed in Chapter 9, “64-bit Support.”

Your compiler must have some means of indicating additional directories to be searched for “includes” of the form:

```
#include <filename>
```

Our reference implementation has only been tested on machines with 32-bit pointers and that do not require “far” pointers of any sort. We do not know if it will run successfully on platforms with pointers of other sizes.

The codebase has been successfully compiled with the following compilers:

- Sun C Compiler 5.0, 5.2 and 5.3 on Solaris,
- GNU C 2.91.66 (egcs-1.1.2) compiler on Red Hat Linux,
- GNU C 2.95.2 compiler on Solaris and Windows NT 4.0,
- Microsoft Visual C++ 6.0 Professional on Windows NT 4.0 and Windows 2000.

The only non-ANSI C feature in the KVM source code base is its use of 64-bit integer arithmetic.

1. Note that in the Java programming language, the type `long` is always 64 bits. Table 3-1 assumes that, as in most current C implementations, the type `long` represents a 32-bit quantity. This document uses the phrase “The Java type `long`” to refer to the 64-bit meaning.

Directory Structure



4.1 Overview

Unzip the release package into any directory of your choice. This will create a directory with the following subdirectories:

- api
- bin
- build
- docs
- jam
- kvm
- samples
- tools

The contents of these directories are detailed in TABLE 4-1.

TABLE 4-1 Distribution directories

Subdirectory	Description
api	Contains the Java library source code that is provided with the release.
bin	Contains all the binary executables and compiled Java library classes.
build	Contains makefiles for building the KVM.
doc	Contains documentation.
jam	Contains the source code of the optional Java Application Manager (JAM) component that is provided with the KVM.

TABLE 4-1 Distribution directories

Subdirectory	Description
kvm	Contains the source code of the KVM.
samples	Contains the source code and icons of a number of sample applications.
tools	Contains the source code and icons of a number of tools (JavaCodeCompact, preverifier, KDWP Debug Proxy (kdp), Palm tools) that are provided with the release.

4.2 Directory `kvm/VmCommon`

All common, platform-independent source code of KVM is located in the directory `kvm/VmCommon/src/`. All common include files are in the directory `kvm/VmCommon/h/`.

Port specific source and include files should go into the directories `kvm/VmPort/src/` and `kvm/VmPort/h/`, where *Port* is replaced by the name of your platform (e.g., `kvm/VmWin`, `kvm/VmPilot`, `kvm/VmUnix`.)

Some ports may choose to create a `kvm/VmPort/build/` subdirectory which holds files that are part of the build process, but are not part of the source code *per se*.

TABLE 4-2 gives an overview of the KVM source code files contained in `kvm/VmCommon/src/` and `kvm/VmCommon/h/`.

TABLE 4-2 Files in `VmCommon`

File	Description
<code>StartJVM.c</code>	Virtual machine startup and command line argument reading.
<code>bytecodes.c</code>	The definition of Java bytecodes for the redesigned bytecode interpreter (since KVM 1.0.2).
<code>cache.h</code> <code>cache.c</code>	Inline caching operations for speeding up method lookup and for supporting “fast” bytecodes.
<code>class.h</code> <code>class.c</code>	Internal runtime data structures and operations for representing Java classes.
<code>events.h</code> <code>events.c</code>	Event system implementation.
<code>execute.h</code> <code>execute.c</code>	Interpreter execution macros and operations needed by the redesigned bytecode interpreter (since KVM 1.0.2).

TABLE 4-2 Files in VmCommon

File	Description
fields.h fields.c	Internal runtime data structures and operations for representing fields and methods.
frame.h frame.c	Stack frame and exception handling operations.
garbage.h garbage.c collector.c collectorDebug.c	Garbage collector and memory management.
global.h global.c	Miscellaneous global variables and definitions.
hashtable.h hashtable.c	Hashtable implementation that is used internally by the virtual machine.
interpret.h interpret.c	Bytecode interpreter. Note that starting from KVM 1.0.2 the actual interpreter code and bytecode definitions are located in other files (bytecodes.c, execute.h, execute.c).
loader.h loader.c	Class loader.
log.h log.c	Logging/diagnostic operations for debugging and profiling.
long.h	Special macros to handle 64-bit operations in a portable fashion.
main.h	Compilation options and system-wide default settings.
messages.h	Error and warning messages.
native.h native.c nativeCore.c	Native function table operations and core native library functions.
pool.h pool.c	Runtime data structures and operations for representing constant pools.
profiling.h profiling.c	Data declarations and operations for profiling virtual machine execution.
property.h property.c	Operations for accessing Java system properties.
rom.h	Macros needed by the ROMizer (JavaCodeCompact tool).
runtime.h	Function templates for certain machine-specific operations that need to be defined for each KVM port.

TABLE 4-2 Files in VmCommon

File	Description
stackmap.c	Stackmap operations that are used for supporting exact garbage collection.
thread.h thread.c	Internal runtime data structures and operations for multithreading and Java thread management.
verifier.h verifier.c	Classfile verifier (see Chapter 12 for details).

4.3 Directory `kvm/VmExtra`

The directory `kvm/VmExtra/` contains additional components that are potentially useful to a large number of ports. These files include an implementation of the most commonly needed networking protocols for Windows/Unix, a file interface for supporting class loading on those target platforms that have a regular file system, and a JAR file reader/inflater. This directory also contains the implementation of the Java-level debugger and the KDWP (KVM Debug Wire Protocol) interface.

In addition, the directory defines some optional macros for asynchronous event handling, and defines the virtual machine startup operations that are needed on non-embedded, command line based target platforms such as Windows and Solaris.

A description of the `VmExtra` files is provided in TABLE 4-3.

TABLE 4-3 Files in `VmExtra`

File	Description
<code>async.h</code>	Macros for supporting asynchronous I/O (see Section 10.3, “Asynchronous native methods” and Section 11.1.4, “Asynchronous notification”).
<code>loaderFile.c</code>	Low-level binding between the file system, class loader and JAR reader for those platforms that have a “real” file system.
<code>main.c</code>	Default main program for those platforms that have a file system and support VM startup from a command line.
<code>jar.h</code> <code>inflate.h</code> <code>inflateint.h</code> <code>inflatetables.h</code> <code>jar.c</code> <code>inflate.c</code>	Jar file reader and inflater (decompressor).
<code>commProtocol.h</code> <code>commProtocol.c</code> <code>socketProtocol.h</code> <code>socketProtocol.c</code> <code>datagramProtocol.h</code> <code>datagramProtocol.c</code>	Implementation of the most commonly used network protocols for Windows/Unix (serial communication ports, sockets, server sockets, datagrams).
<code>resource.c</code>	Implementation of a stream-based protocol for reading external resources.

TABLE 4-3 Files in VmExtra

File	Description
debugger.h debugger.c debuggerCommands.h debuggerStreams.h debuggerInputStream.c debuggerOutputStream.c debuggerSocketIO.c	Implementation of the Java-level debugger and the KDWP (KVM Debug Wire Protocol) interface.
fakeStaticMemory.c	Memory management definitions that allow KVM emulate the special USESTATIC mode on Windows/Unix for debugging purposes.
nativeSpotlet.h nativeSpotlet.c	Defines the low-level operations for event handling, and the routines to intercept events from the host operating system and invoke the registered handler methods for events.

Required Port-Specific Files and Functions

This section describes those files and functions that must be defined for each port.

5.1 File `machine_md.h`

Every KVM port must provide a file named `VmPort/h/machine_md.h`. The purpose of this file is to override the default compile time definitions and declarations provided in `VmCommon/h/main.h`, and supply any additional definitions and declarations that your specific platform might need. See Chapter 6, “Compilation Flags, Definitions and Macros” for a list of the definitions and declarations that your port will often need to override.

All port-specific declarations, function prototypes, typedef statements, `#include` statements, and `#define` statements must appear either in this `machine_md.h`, in a file included directly or indirectly by `machine_md.h`, in some file automatically included by your development environment,¹ or via compiler switches.²

Port-specific functions can appear in any machine-specific file. Unless otherwise stated, any required port-specific function can also be defined as a macro, provided that its implementation is careful to ensure that each argument is evaluated exactly once.

1. Metrowerks CodeWarrior, for example, allows the user to create a *prefix file*.

2. Some compilers allow you to add the switch `-Dname=value`, which is equivalent to putting `#define name value` at the start of the file.

5.2 File `main.c`

You will generally need to provide a new version of `main.c` that is suitable for your target platform. The default implementation provided in directory `VmExtra/src/main.c` can be used as a starting point for platform-specific implementations. Refer to Chapter 7, “Virtual Machine Startup,” for further information.

5.3 Runtime functions that require porting efforts

Each port must define the functions given below (see `VmCommon/src/runtime.h`). They may be defined as either macros or as C code. Traditionally, the C code is placed in a file named `VmPort/src/runtime_md.c`

- `void AlertUser(const char* message)`
Alert the user that something serious has happened. This function call usually precedes a fatal error.
- `cell *allocateHeap(long *sizeptr, void **realresultptr)`
Create a heap whose size (in bytes) is approximately the long value `*sizeptr`. The heap must begin at an address that is a multiple of 4. The address of the heap is returned as the value of this function. The actual size of the heap (in bytes) is returned in `*sizeptr`. The value placed into `*realresultptr` is used as the argument to `freeHeap` when freeing the heap.
For most ports, `*realresultptr` will be set to the actual value returned by the native space allocation function. If this value is not a multiple of 4, it is rounded up to the next multiple of 4, and `*sizeptr` is decreased by 4.
- `void freeHeap(void *heapPtr)`
Free the heap space that was allocated using `allocateHeap`. See above for the meaning of the `heapPtr` argument.
- `GetAndStoreNextKVMEvent(bool_t forever, ulong64 waitUntil)`
This function serves as an interface between the event handling capabilities of the virtual machine and the host operating system. See Chapter 11 for details.
- `void InitializeVM()`
Initialize the virtual machine in whatever way is necessary. On many of the current ports, this is a macro that does nothing.
- `void InitializeNativeCode()`
Initialize the native code in whatever way is necessary. Ports can use this function (for example) to initialize the window system and to perform other native-code specific initialization.

- `void InitializeClassLoading()`
Initialize the class loader in whatever way is necessary. Ports can use this function (for example) to perform certain file/storage system initialization operations.
- `void FinalizeVM()`
Perform any cleanup necessary before shutting down the virtual machine.
- `void FinalizeNativeCode()`
Perform any clean up necessary to clean up after the native functions. Many ports use this function to shut down the window system.
- `void FinalizeClassLoading()`
Perform any cleanup necessary before shutting the class loader. Ports can use this function (for example) to perform certain file/storage system finalization operations.
- `ulong64 CurrentTime_md(void)`
Return the time, in milliseconds, since January 1, 1970 UTC. On devices that do not support the concept of time zone, it is acceptable to return the time, in milliseconds, since January 1, 1970 of the current time zone.

The functions `InitializeNativeCode()` and `InitializeVM()` are called, in that order, before any global variables have been set and before the memory-management system has been initialized.

The function `FinalizeVM()` is called just before `FinalizeNativeCode()`. On those ports that have enabled profiling, the profiling information is printed out between the calls to these two functions. This allows the profiler to find out information about the window system, if necessary, and to use the window system for creating its output.

Note – If you want to use the KVM for running additional libraries such as those defined by the *Mobile Information Device Profile (MIDP)* or *PDA Profile*, additional porting work will be necessary to port the native functions required by those libraries.

Asynchronous native functions. If your port supports the use of asynchronous native methods, there are additional, port-specific functions that you must define:

```
yield_md()
CallAsyncNativeFunction_md()
enterSystemCriticalSection()
exitSystemCriticalSection()
```

These functions are described in §10.3.

5.4 Required C library functions

The KVM uses the following C library functions:

- String manipulation: `strcat`, `strchr`, `strcmp`, `strcpy`, `strncpy`, `strlen`
- Moving memory: `memcpy`, `memmove`, `memset`, `memcmp`
- Printing: `atoi`, `sprintf`, `fprintf`, `putchar`
- Exception handling: `setjmp`, `longjmp` (not absolutely necessary)

If your development environment does not supply definitions for these functions, you must either define them yourself, or use macros to map these names onto equivalent functions recognized by your development environment.¹

The function `memmove` must be able to handle situations in which the source and destination overlap. The function `memcpy` is used only in those cases in which the source and destination are known not to overlap.

The functions `fprintf` and `sprintf` use the following formats:

`%s`, `%d`, `%o`, `%x`, `%ld`, `%lo`, `%lx`, `%%`

These formats never have options or flags.

There are no calls directly to `printf`.

Note – The components included in directory `VmExtra`, the machine-specific ports provided with this release, and the optional Java Application Manager (JAM) component may need additional native functions not listed above.

1. Be aware that the order of arguments may be different on different platforms. For example, the function `memset` takes arguments `memset(location, value, count)`. The corresponding Palm OS function is `MemSet(location, count, value)`.

Compilation Flags, Definitions and Macros

This section lists various C preprocessor flags, definitions and macros that are defined `VmCommon/h/main.h`. Understanding the meaning of these flags helps you in porting efforts, so please read the documentation below and in file `VmCommon/h/main.h`.

Note – Rather than changing the values provided in `VmCommon/h/main.h`, these values should be preferably be overridden in your port-specific `machine_md.h` file.

Also note that in our reference implementation, many of these flags are commonly overridden from makefiles.

For each definition, we give a brief summary and its default definition. These flags and macros are documented also in `VmCommon/h/main.h`.

6.1 General compilation options

The following definitions control the general platform-dependent compiler options that you must set before starting your porting efforts. Incorrect settings typically cause the virtual machine to malfunction.

```
#define COMPILER_SUPPORTS_LONG 1
```

Turn this flag on if your compiler has support for long (64 bit) integers.

```
#define NEED_LONG_ALIGNMENT 0
```

Instructs the KVM to know that your host operating system and compiler generally assume all 64-bit integers to be aligned on eight-byte boundaries.

```
#define NEED_DOUBLE_ALIGNMENT 0
```

Instructs the KVM to know that your host operating system and compiler generally assume all double floating point numbers to be aligned on eight-byte boundaries (this flag is meaningful only if floating point support is turned on.)

Additional notes. The compiler generates better code if it knows the “endianness” of your machine. You should set one of the following two variables to “1” in your machine-specific header file.

```
#define BIG_ENDIAN 0
#define LITTLE_ENDIAN 0
```

It is necessary to set one of these “endian” variables to “1” if you reset `COMPILER_SUPPORTS_LONG` to zero. (See Chapter 9 for more details.)

Also note that if your compiler supports 64-bit integer arithmetic and you have set the flag

```
#define COMPILER_SUPPORTS_LONG 1
```

you should supply definitions for the types `long64` and `ulong64`. If your compiler does not support 64-bit integers (or you have set the flag to 0 for some other reason), structure definitions of these two types are created for you automatically. (See Chapter 9.)

6.2 General system configuration options

The following definitions allow you to control which components and features to include in your port.

```
#define IMPLEMENTS_FLOAT 0
```

Turns floating point support in KVM on or off. Should be off in those implementations that are compliant with CLDC Specification 1.0.

```
#define PATH_SEPARATOR ':'
```

Path separator character used in `CLASSPATH`. This definition is meaningful only when utilizing the default class loader for command line based systems.

```
#define ROMIZING 1
```

Turns class prelinking/preloading (JavaCodeCompact) support on or off. If this option is turned on, KVM prelinks all the system classes directly in the virtual machine, speeding up application startup considerably. Refer to Chapter 13 for details.

```
#define USE_JAM 0
```

Includes or excludes the optional Java Application Manager (JAM) component in the virtual machine. Refer to Chapter 14 for details.

```
#define ASYNCHRONOUS_NATIVE_FUNCTIONS 0
```

Instructs the KVM to use optional asynchronous native functions. Refer to Section 10.3, “Asynchronous native methods” and Chapter 11 for details.

6.3 Palm-specific system configuration options

The following definitions allow you to control certain Palm-specific system configuration options. All these features were originally designed for Palm OS version of KVM, but they may be useful also for other ports.

```
#define USESTATIC 0
```

Instructs the KVM to use a Palm-specific optimization in which certain immutable runtime data structures are moved from “dynamic RAM” to “storage RAM” to conserve Java heap space. A fake implementation of this mechanism is available also for the Windows and Solaris versions of KVM (for debugging purposes.)

```
#define CHUNKY_HEAP 0
```

Instructs the KVM to use an optimization which allows the KVM to allocate the Java heap in multiple chunks or segments. This makes it possible for the virtual machine to allocate more heap space on certain platforms such as Palm OS.

```
#define RELOCATABLE_ROM 0
```

Instructs the KVM to use an optimization in which the prelinked system classes are stored using a relocatable (movable) representation. This allows romized (JavaCodeCompacted) system classes to be stored in devices such as Palm OS.

6.4 Memory allocation settings

The following definitions affect the amount of memory KVM allocates.

```
#define DEFAULTHEAPSIZE 65024 /* 0xFE00 */
```

The Java heap size that KVM allocates upon virtual machine startup. This value is commonly overridden from makefiles. Note that, starting from KVM 1.0.3, it is possible to override the heap size value from the command line (in those ports that support command line operation.)

```
#define INLINECACHE_SIZE 100
```

The size of a special inline cache area that KVM reserves upon virtual machine startup if the `ENABLEFASTBYTECODES` option is turned on. The inline caching mechanism speeds up method lookups in the KVM by utilizing a technique popularized by Deutsch & Schiffman in the early 1980s. The size here is expressed as a number of inline cache entries (each entry requires 12-16 bytes depending on your target platform.)

```
#define STACKCHUNK_SIZE 128
```

The execution stacks of Java threads inside the KVM grow and shrink automatically as necessary. This value defines the default size of a new stack frame chunk when a new stack chunk needs to be allocated. Reducing the default stack chunk size will make the creation of new Java threads less expensive, but will slow down the execution of the VM when running programs that require a lot of stack space (that is, programs that have a lot of nested method calls.)

```
#define STRINGBUFFER_SIZE 512
```

The size (in bytes) of a statically allocated area that the virtual machine uses internally in various string operations.

Note – As a general principle, KVM allocates all the memory it needs upon virtual machine startup. At runtime, all the memory is allocated inside the preallocated areas. Of course, the situation may change if the virtual machine calls host-system specific native functions (such as graphics functions) that perform dynamic memory allocation outside the Java heap.

6.5 Garbage collection options

The following option turns on compacting garbage collection. Note that currently compaction cannot be used on those platforms that have a segmented (non-contiguous) memory architecture.

```
#define ENABLE_HEAP_COMPACTION 1
```

The following option, if set to a non-zero value, causes a garbage collection to occur on every allocation. This makes it easier to find garbage collection problems. Since this option makes the virtual machine run extremely slowly, the option should be turned off in production builds.

```
#define EXCESSIVE_GARBAGE_COLLECTION 0
```

6.6 Class loading options

Some KVM ports may want to forbid any new classes from being loaded into any system package. The following macro defines whether a package name is one of these restricted packages. By default, the system prevents dynamic class loading to `java.*` and `javax.*` packages.

```
#define IS_RESTRICTED_PACKAGE_NAME(name) \
((strcmp(name, "java.", 5) == 0) || \
 (strcmp(name, "javax.", 6) == 0))
```

6.7 Interpreter execution options (KVM 1.0)

The following macros allow you to turn on and off certain features controlling interpreter execution. The default values for a production release are shown below.

```
#define ENABLEFASTBYTECODES 0
```

Turns runtime bytecode replacement and method inline caching on or off. This option improves the performance of the virtual machine by about 10-20%, but increases the size of the virtual machine by a few kilobytes. Note that bytecode replacement cannot be performed on those target platforms in which bytecodes are stored in non-volatile memory such as ROM.

```
#define VERIFYCONSTANTPOOLINTEGRITY 1
```

Instructs the virtual machine to verify the types of constant pool entries at runtime when performing constant pool lookups. Reduces runtime performance slightly, but is generally recommended to be kept on for safety and security reasons.

Additional definitions and interpreter macros:

```
#define BASETIMESLICE
```

The value of this variable determines the basic frequency (as a number of bytecodes executed) in which the virtual machine performs thread switching, event notification and other periodically needed operations. A smaller number reduces event handling and thread switching latency, but causes the interpreter to run more slowly.

```
#define DOUBLE_REMAINDER(x, y) fmod(x,y)
```

A compiler macro, defined in `interpret.h`, that is used to find the modulus of two floating point numbers.

```
#define SLEEP_UNTIL(wakeupTime)
```

This macro makes the virtual machine sleep until the current time (as indicated by the return value of the function `CurrentTime_md()`) is greater than or equal to the wakeup time. The default implementation of `SLEEP_UNTIL` is a busy loop. Most ports should usually provide a more efficient implementation for battery conservation reasons. Refer to Section 11.4, “Battery power conservation” for further details.

6.8 Interpreter execution options (KVM 1.0.2 and later)

Since the release 1.0.2, KVM has an interpreter design that gives up to 15-30% better performance than KVM 1.0 without any loss of ANSI C portability. The actual performance improvement percentage depends on the target platform and the capabilities of the C compiler that is used for compiling the KVM. The performance improvement is the result of the following four techniques that can be used independently of each other:

- Restructuring the interpreter code so that virtual machine registers will be placed into local C variables when the interpreter is running.
- Splitting uncommonly used Java bytecodes into a separate interpreter loop subroutine. This allows the C compiler to do a better job in optimizing the code for more frequently used bytecodes.

- Moving the test for Java thread rescheduling from the top of the interpreter loop to branch bytecodes. This reduces the overhead of the timeslice counter that is used for controlling thread switching.
- Padding out the bytecode space in order to allow the C compiler to produce better code for the main switch statement of the interpreter.

These techniques do not depend on any compiler-specific features, and are therefore portable across a wide variety of C compilers. Each of the techniques and the corresponding macros are discussed in more detail below.

6.8.1 Copying the virtual machine registers to local variables

The virtual machine registers of the KVM (`ip`, `sp`, `lp`, `fp`, `cp`) are accessed very frequently when bytecodes are being executed. In KVM 1.0, all these virtual machine registers are defined as global C variables. Starting from KVM 1.0.2, these registers are still principally defined as global variables, but if the `LOCALVMREGISTERS` option is on, they are copied to local variables when the interpreter is executing. A good C compiler will then optimize the interpreter loop so that these local variables are put into machine registers for substantially faster execution.

```
#define LOCALVMREGISTERS 1
```

Turns the localization of virtual machine registers on or off.

```
#define IPISLOCAL 1
#define SPISLOCAL 1
#define LPISLOCAL 0
#define FPISLOCAL 0
#define CPISLOCAL 0
```

These macros allow you to control specifically which of the virtual machine registers should be used locally by the interpreter loop. These macros have been added to provide better control over register allocation, as many resource-constrained platforms may not have many physical hardware registers available.

The optimal selection of these options for a specific platform will require careful examination of the machine code produced by the compiler, along with a good deal of experimentation. By default, `ip` (instruction pointer), and `sp` (stack pointer) are allocated locally, while `lp` (locals pointer), `fp` (frame pointer) and `cp` (constant pool pointer) are kept in global variables.

Note – If you use the `LOCALVMREGISTERS` option and you want to make further changes to the code implementing Java bytecodes, the single most important thing to remember is to make sure that the local copies of the virtual machine registers are copied back to their global variables before calling functions in the virtual machine that expects them to be in their global variables. Failure to do so will lead to obscure bugs. The virtual machine registers can be saved to their global variables by using the macro `VMSAVE`. They are restored back to their local variables by using the macro `VMRESTORE`. For instance the `RETURN` bytecodes may need to call `monitorExit()`, and to do this the call must be done as follows:

```
VMSAVE
result = monitorExit(...);
VMRESTORE
```

6.8.2 Splitting uncommon bytecodes into a separate subroutine

The KVM 1.0 interpreter had the code for all the Java bytecodes in a single large switch statement. However, a majority of Java bytecodes are executed very rarely. If the code for the more frequently and less frequently used bytecodes is placed in separate routines, the C compiler can often do a better job optimizing the resulting smaller interpreter loops. This also helps the compiler find hardware registers for the virtual machine registers more easily when the `LOCALVMREGISTERS` option is in use.

```
#define SPLITINFREQUENTBYTECODES 1
```

Turning this option on allows the C compiler to generate separate interpreter loops for the frequently and infrequently used bytecodes.

Note that the code to process the bytecodes is now contained in a file called `bytecodes.c`. The code for all the bytecodes is kept here and is selectively compiled by utilizing a number of internal macro definitions (`STANDARDBYTECODES`, `INFREQUENTSTANDARDBYTECODES`, `FLOATBYTECODES` and `FASTBYTECODES`).

The code in `bytecodes.c` is executed from another new file called `execute.c`. If the `SPLITINFREQUENTBYTECODES` option is enabled, the file `bytecodes.c` is included twice into `execute.c`: once for the routine called `SlowInterpret()` and once for the routine `Interpret()`. The four macros mentioned above are used to control the expansion of the appropriate bytecodes into the correct subroutines.

6.8.3 Moving the test for thread rescheduling to branchpoints

The old KVM 1.0 interpreter tested for the need to reschedule (switch threads) before the execution of each bytecode. The performance of the interpreter was improved by about 5% by changing the location of this test so that the test is performed only after every branch, goto, call and return instruction.

Thread scheduling in the old interpreter took place when a certain number of bytecodes had been executed. This number was, by default, 100 times the priority of the thread. In the new interpreter (since KVM 1.0.2), thread rescheduling occurs by default when 100 times the number of branch, call, or return bytecodes have been executed.

```
#define RESCHEDULEATBRANCH 1
```

Turning this option on changes the thread switching mechanism so that tests for thread switching are moved to branchpoints. Note that enabling this option affects the value of the `BASETIMESLICE` macro inherited from KVM 1.0. When this option is off, thread scheduling operates as in KVM 1.0.

6.8.4 Padding out the bytecode space

The Java Virtual Machine Specification defines 200 standard bytecodes, plus additionally reserves four other bytecodes for other use. However, many C compilers produce better code when the size of the bytecode (switch) table is exactly 256.

```
#define PADTABLE 0
```

Turning this option on will pad the interpreter switch tables so that the number of instructions is 256. This will increase the size of the virtual machine, but allows the interpreter to run faster on some platforms.

6.9 Java-level debugging options

The KVM 1.0.2 release introduced a new Java-level debugger interface that allows the KVM to be plugged into third party Java debugger environments and integrated development environments (IDEs). The macros in this subsection are related to the Java-level debugger options.

Note – It is important to notice that there is a fundamental difference between the debugging facilities intended for *Java-level debugging* and *VM-level debugging*.

Java-level debugging facilities are related to the debugging of the Java programs that the KVM executes. *VM-level debugging* facilities are used for debugging the KVM itself at the native (C) code level.

```
#define ENABLE_JAVA_DEBUGGER 1
```

Includes a large amount of debugger support code that is needed for plugging KVM into a third-party Java debugger or integrated development environment such as Forte or Borland JBuilder.

More information about the Java-level debugger facilities and the KDWP interface is provided in Chapter 15, “Java-Level Debugging Support (KDWP).”

6.10 VM-level debugging and tracing options

KVM provides a large number of debugging and tracing facilities that can be used for inspecting the behavior of the KVM itself at the native (C) code level. These facilities can be extremely helpful during porting efforts.

All the VM-level debugging and tracing options should be turned off in a production release.

6.10.1 Including and excluding debugging code

```
#define INCLUDEDDEBUGCODE 0
```

Includes a large amount of debugging and logging code that is useful when porting the virtual machine onto a new platform. This option should be turned off in production builds.

```
#define ENABLEPROFILING 0
```

Turns on or off certain profiling features that allow you to monitor virtual machine execution and get execution statistics. Turning this option on slows down the virtual machine execution speed considerably. This option should be turned off in production builds.

6.10.2 Tracing options

In KVM 1.0, all the tracing options were compilation flags that could be changed only by recompiling the virtual machine. In KVM 1.0.2, all these tracing options were changed into global variables that can be controlled from the command line. This makes it much easier to turn individual tracing options on and off. These global variables (and command line switches) are available only if the virtual machine has been compiled with the `INCLUDEDEBUGCODE` mode turned on.

<code>-traceallocation</code>	trace memory allocation
<code>-tracedebugger</code>	trace the debugging interface (since KVM 1.0.3)
<code>-tracegc</code>	trace garbage collection
<code>-tracegcverbose</code>	trace garbage collection, more verbose
<code>-traceclassloading</code>	trace class loading
<code>-traceclassloadingverbose</code>	trace class loading, more verbose
<code>-traceverifier</code>	trace class file verifier
<code>-tracestackmaps</code>	trace the behavior of stack maps
<code>-tracebytecodes</code>	trace bytecode execution
<code>-tracemethods</code>	trace method calls
<code>-tracemethodsverbose</code>	trace method calls, more verbose
<code>-traceframes</code>	trace stack frames
<code>-tracestackchunks</code>	trace the allocation of new stack chunks
<code>-traceexceptions</code>	trace exception handling
<code>-traceevents</code>	trace the behavior of the event system
<code>-tracethreading</code>	trace the behavior of the multithreading system
<code>-tracemonitors</code>	trace the behavior of monitor objects
<code>-tracenetworking</code>	trace the network access
<code>-traceall</code>	activates all the tracing options above simultaneously

If your target platform does not support command line operation, you can control these options directly by changing their default values in file `VmCommon/src/global.c`, or by defining a graphical user interface that sets and resets these options.

Additionally, you can control whether the tracing messages printed out are terse or more verbose by modifying the following option:

```
#define TERSE_MESSAGES 0
```

KVM also contains a stack trace printing facility that can be turned on to help debugging of exceptions and errors in more detail (at the cost of some additional memory footprint). By default, this mode is turned on automatically when the `INCLUDEDEBUGCODE` flag is turned on.

```
#define PRINT_BACKTRACE 0
```

6.11 Error handling macros

The interpreter uses code of the form shown in Figure 6-1.

If there is a call to the macro `ERROR_THROW(int)`, anywhere inside the “normal code,” the VM jumps immediately to error handling code. Uses of this macro can be nested, either lexically or dynamically. The `ERROR_THROW` jumps to the innermost `ERROR_CATCH` error handling code.

```
ERROR_TRY {  
    normal code  
} ERROR_CATCH (error) {  
    error handling code  
} ERROR_END_CATCH  
    always continue here
```

FIGURE 6-1 Error handling

By default, this behavior is emulated using `set jmp` and `long jmp`. However, platforms (such as PalmOS) that already provide a similar mechanism should use the native mechanism.

6.12 Miscellaneous macros and options

```
#define UNUSEDPARAMETER(var)
```

Some functions in the reference implementation take arguments that they do not use. Some compilers issue warnings; others do not. For those compilers that do issue warnings, they differ in how you indicate that the non-use of the variable is intentional and that you do not wish to get a warning. This macro should do whatever is necessary to get your compiler to remain quiet.

6.13 Overriding the compilation flags and other options from makefiles

The following parameters are commonly used when using `gnumake` to build the KVM.

```
gnumake ROMIZING=true
```

Build the KVM with romizing enabled, that is, link all the system classes statically into the KVM executable.

```
gnumake DEBUG=true
```

Build the KVM with the Java-level debugger enabled.

```
gnumake USE_JAM=true
```

Build the KVM with the Java Application Manager (JAM) enabled.

```
gnumake GCC=true
```

Use GNU C compiler instead of the standard Sun compiler (on Solaris.)

Virtual Machine Startup

Virtual machine startup practices can vary significantly in different KVM ports. By default, KVM supports regular command line based Java virtual machine startup, but the virtual machine can easily be modified for those environments in which command line based startup is not desired.

7.1 Command line startup

This subsection describes the virtual machine startup conventions when launching KVM from a command line.

The file `VmExtra/src/main.c` provides a default implementation of `main()`. The virtual machine is called from the command line as follows:

```
kvm [option]* className [arg]*
```

where each *option* is one of

```
-version  
-classpath <list of directories>  
-heapsize <heap size parameter>
```

The required *className* argument specifies the class whose method `static main(String argv[])` is to be called. All arguments beyond the class name are uninterpreted strings that are made into a single `String[]` object and passed as the single argument to the `main` method.

The `-classpath` option allows the user to define the directories from which the KVM reads the class files. The parameter `<list of directories>` is a single string in which the directories are separated by the `PATH_SEPARATOR` character. The value of the `PATH_SEPARATOR` character is typically `'/'` on Windows platforms, and `':'` on Unix platforms.

The `-heapsize` option (introduced in KVM 1.0.3) allows the user to manually set the Java heap size that KVM allocates upon virtual machine startup. The heap size can range from 32 kilobytes to 64 megabytes. The heap size can be specified either in bytes (e.g., 32768), kilobytes (e.g., 32k, 32kB, 32K or 32KB), or megabytes (e.g., 1m, 1M, 1MB or 1MB). Note that when the heap size is defined in bytes, the KVM automatically rounds up the heap size number to the next number that is divisible by four.

Additionally, if the virtual machine has been compiled with the `INCLUDEDEBUGCODE` mode turned on, the following tracing options are available:

<code>-traceallocation</code>	trace memory allocation
<code>-tracedebugger</code>	trace the debugging interface (new in KVM 1.0.3)
<code>-tracegc</code>	trace garbage collection
<code>-tracegcverbose</code>	trace garbage collection, more verbose
<code>-traceclassloading</code>	trace class loading
<code>-traceclassloadingverbose</code>	trace class loading, more verbose
<code>-traceverifier</code>	trace class file verifier
<code>-tracestackmaps</code>	trace the behavior of stack maps
<code>-tracebytecodes</code>	trace bytecode execution
<code>-tracemethods</code>	trace method calls
<code>-tracemethodsverbose</code>	trace method calls, more verbose
<code>-traceframes</code>	trace stack frames
<code>-tracestackchunks</code>	trace the allocation of new stack chunks
<code>-traceexceptions</code>	trace exception handling
<code>-traceevents</code>	trace the behavior of the event system
<code>-tracethreading</code>	trace the behavior of the multithreading system
<code>-tracemonitors</code>	trace the behavior of monitor objects
<code>-tracenetworking</code>	trace the network access
<code>-traceall</code>	activates all the tracing options above simultaneously

When the Java-level debugging interface is in use, additional command line options are available to control the debugger. Refer to Chapter 15 for details.

The default implementation of `main(int argc, char **argv)` calls the function `StartJVM()` with an `argv` in which all of the options have been removed and an `argc` that has been decremented appropriately.

7.2 Alternative VM startup strategies

If your implementation does not start the virtual machine from a command line (for example, if you use a graphical environment for application launching), you must arrange your code to call `StartJVM()` with the appropriate arguments.

7.3 Using a JAM (Java Application Manager)

Many KVM ports run on resource-constrained devices which lack many features commonly available in desktop operating systems, e.g., a command line language, graphical file manager, or even a file system. To facilitate the porting of KVM to such platforms, KVM provides a sample implementation of a facility called JAM (Java Application Manager).

At the compilation level, JAM can be turned on or off by using the flag

```
#define USE_JAM 1
```

When building the KVM using `gnumake`, the following command automatically builds the system with the JAM enabled:

```
gnumake USE_JAM=true
```

If JAM is compiled into the KVM, it must be activated with the `-jam` command line flag.

The JAM implementation assumes that applications are available for downloading as JAR files by using the HTTP protocol. The JAM reads the contents of the JAR file and an associated descriptor file via HTTP, and launches KVM with the main class as a parameter.

Since the JAM serves as an interface between the host operating system and the virtual machine, it can be used, e.g., as a starting point for a device-specific graphical Java application management and launching environment (“microbrowser”), or as a test harness for virtual machine testing. The JAM reference implementation provides a special “`-repeat`” mode that allows the JAM to run a large number of Java applications (e.g., test cases) without having to restart the virtual machine every time.

Refer to Chapter 14, “Java Application Manager (JAM),” for further information on the JAM.

Class Loading, JAR Files, and Inflation

The KVM source code includes an implementation for reading Java class files from regular files/directories, as well as (possibly compressed) JAR files.

If you need to provide an alternative method for loading class files, you must define your own class loading mechanism. The default implementation in `VmExtra/src/loaderFile.c` can be used as a starting point for platform-specific implementations.

The KVM code to read JAR files can also be used independently of reading class files. Applications that need to make their own use of JAR files can use these functions. In addition, the function that decompresses compressed JAR entries (a process called “inflation”), can also be used to decompress other information. For example, the PNG image format uses the same compression and decompression algorithms.

The code to read JAR files and to inflate compressed files has been written so that it can also be used in programs other than the KVM. By setting appropriate preprocessor symbols, these functions can use `malloc()` and `free()` rather than the KVM heap.

8.1 Generic class file loading

You must define the C structure `filePointerStruct`. The generic code uses the definitions

```
struct filePointerStruct;  
typedef struct filePointerStruct *FILEPOINTER;
```

without knowing anything about the fields of this structure.

You must also define the following functions:

- `void InitializeClassLoading()`
The code typically initializes the variable `ClassPathTable` and any other variables needed for file loading upon virtual machine startup. Keep in mind that the value in `ClassPathTable` is usually a root for garbage collection, and must either be `NULL` or be an object allocated from the heap.
- The C preprocessor constant `PATH_SEPARATOR` indicates the character that separates directories in the class path. Its default value is `'.'`. If you are using Windows or a similar implementation, you will need to change this value to `'\'`.
- `void FinalizeClassLoading()`
This function is the opposite of `initializeClassLoading()`. This function performs the class loader finalization operations that are necessary when the virtual machine shuts down. Actual implementation will vary substantially depending on the target architecture.
- `FILEPOINTER openClassfile(const char *className)`
Open the class file containing the class whose name is `className`. The variable `className` is a fully qualified class name that uses slashes (`'/'`) as the package separator.
- `void closeClassfile(FILEPOINTER ClassFile)`
Close the indicated class file. Close any system resources (such as file handles or database records) associated with the class file.
- `void loadByteNoEOFCheck(FILEPOINTER ClassFile)`
Load the next byte if it is a JAR file, or load the next character and return it, or EOF (-1) if end of file was reached.
- `unsigned char loadByte(FILEPOINTER ClassFile)`
`unsigned short loadShort(FILEPOINTER ClassFile)`
`unsigned long loadCell(FILEPOINTER ClassFile)`
Read the next one, two, or four bytes from the class file, and return the result as an unsigned 8-bit, unsigned 16-bit, or unsigned 32-bit value. 16- and 32-bit quantities in Java class files are always in big-endian format.
- `void loadBytes(FILEPOINTER ClassFile, char *buffer, int len)`
Load the next `len` bytes from the class file into the indicated buffer.
- `void skipBytes(FILEPOINTER ClassFile, unsigned int len)`
Skip the next `len` bytes in the class file.

The class file structure returned by `openClassFile` must be an object allocated from the Java heap.

8.2 JAR file reader

CLDC-compliant KVM implementations are required to be able to read class files from compressed JAR files. The location of the JAR file(s) is specified in an implementation-dependent manner.

Functions are provided in `jar.c` for reading entries a JAR file. If the preprocessor symbol `JAR_FILE_USES_STDIO` is non-zero, then these functions use C standard I/O routines to read the JAR file. If this preprocessor symbol is set to 0, this indicates that JAR files are in memory.

You can use the JAR file reader outside of the KVM by setting the C preprocessor variable `COMPILING_FOR_KVM` to 0. In this case, the code uses `malloc()` and `free()` to allocate memory.

The JAR file reader uses the inflater, which is discussed in the next section.

8.2.1 Opening a JAR file

Before using a JAR file, you must “open” it using the function

```
bool_t openJARFile(void *nameOrAddress, int length,
                  JAR_INFO entry)
```

The arguments are as follows:

If `JAR_FILE_USES_STDIO` is non-zero, then the first argument is the name of the JAR file and the second argument is ignored.

If `JAR_FILE_USES_STDIO` is zero, then the first argument is a pointer in memory to the beginning of the JAR file, and the second argument is the length, in bytes, of the JAR file.

The third argument is a pointer to a structure of type `struct jarInfoStruct` defined in `jar.h`. This structure is filled with information about the opened JAR file. This function returns `TRUE` if it successfully managed to open the JAR file and parse its directory; it returns `FALSE` otherwise.

8.2.2 Closing a JAR file

If a JAR file has been successfully opened using `openJARFile`, you must close the file when you are done. You must use the function:

```
void closeJARFile(JAR_INFO entry)
```

The last argument is a pointer to the same structure that was filled in by `openJARFile`.

8.2.3 Reading a JAR file entry

To read a specific entry in a JAR file, you use the function

```
void *
loadJARFileEntry(JAR_INFO jarFile,
                  const char *filename,
                  long *length, int extraBytes);
```

The `jarFile` argument is a pointer to the structure filled in by `openJARFile`. The `filename` argument is the null-terminated name of the entry.¹ The `extraBytes` entry indicates that the JAR reader should pad the result with that many extra bytes at the beginning.

If the JAR file reader is successful, it will set the `*length` argument to the length of JAR file entry. This length does *not* include padding inserted because of the `extraBytes` argument. The actual entry (plus padding) is returned as the result of this function.

If the JAR file reader could not find the entry, or if for some reason it was unable to read the entry, this function returns `NULL`.

The result of this function is a heap-allocated object. If this function is called from within the KVM, then you must protect it, as necessary, from garbage collection.

8.2.4 Reading multiple JAR file directory

To read the directory of a JAR file and possibly some of its entries, use the function

```
loadJARFileEntries(JAR_INFO jarFile,
                   JARFileTestFunction testFunction,
                   JARFileRunFunction runFunction,
                   void* info);
```

The `jarFile` argument is a pointer to the structure filled in by `openJARFile`. The `testFunction` and `runFunction` arguments are callback functions whose use is described below. The `info` argument is not used by the jar directory reader, but is passed on an argument to the `testFunction` and `runFunction` callbacks.

The `testFunction` argument is a callback function that is called on each (non-directory) entry in the JAR file. It is called as follows:

1. Note that Jar files always use `'/'` as the directory separator character.

```
typedef bool_t
    (*JARFileTestFunction)(const char *name,
                           int nameLength,
                           int *extraBytes,
                           void *info);
```

The name and nameLength argument specify the name of entry in the JAR file directory. The name argument is *not* null terminated. The value *extraBytes is initially zero, but you can change it to a different value to indicate that the result needs to be padded with extra bytes at the beginning. The info argument is the same as whatever was passed to loadJARFileEntries.

If this function returns TRUE, it indicates that you want to read this entry. If this function returns FALSE, you do not want to read this entry.

For every entry in which testFunction returns TRUE, the jar file reader reads the data and calls the runFunction as follows:

```
typedef void
    (*JARFileRunFunction)(const char *name, int nameLength,
                          void *value, long length, void *info);
```

The name and nameLength arguments are the same as above. The value argument gives the result of reading the JAR file entry. The length argument is the length of the JAR file entry, not including any padding bytes. The info argument is the same as whatever was passed to loadJARFileEntries.

If reading the entry is unsuccessful, then the runFunction is called with the value argument set to NULL.

The value argument is allocated on the heap. If this function is called from within the KVM, then you must protect it, if necessary, from garbage collection.

8.3 Inflation

The inflate function can be used to decompress streams that have been compressed using the so-called deflation algorithm. This is the compression algorithm commonly used in JAR files and in the PNG image format.

You can use the inflate function outside of the KVM by setting the C preprocessor variable COMPILING_FOR_KVM to 0.

The function that inflates JAR file entries can also be used for other purposes. The function is called with the following arguments.

```
typedef int    (*JarGetByteFunctionType)(void *);
```

```

bool_t inflate(void *data, JarGetByteFunctionType
               getByteFunction,
               int compLen,
               unsigned char **outFileHandle,
               int decompLen);

```

This function decompresses a stream of `compLen` bytes into a buffer of `decompLen` bytes. Successive bytes of input are obtained by repeatedly calling

```
getByteFunction(data)
```

This function will be called up to `compLen + INFLATER_EXTRA_BYTES` times, where `INFLATER_EXTRA_BYTES` is defined in `inflate.h` to be the constant 4. Any values returned beyond the first `compLen` calls to the function are immaterial.

The argument `outFileHandle` must be a pointer to a pointer to a buffer of at least `decompLen` characters. If this function is used with the KVM, the buffer must either not be in the heap, or `outFileHandle` must be registered with the garbage collector so that `*outFileHandle` is updated if the buffer is moved.

This function returns `TRUE` if the decompression is successful, and `FALSE` otherwise.

64-bit Support

We do not require your compiler to support 64-bit arithmetic. However, having a 64-bit capable compiler makes porting much easier.

9.1 Setup

Your compiler supports 64-bit integers: You should define the types `long64` and `ulong64` in one of your platform-dependent include files. The meaning of these two types is shown below in Table 9-1.

TABLE 9-1 64-bit types

Type	Description
<code>long64</code>	A signed 64-bit integer.
<code>ulong64</code>	An unsigned 64-bit integer.

You should consider setting one of the two compiler constants `BIG_ENDIAN` or `LITTLE_ENDIAN` to a non-zero value. This is only required if you are using the Java Code Compactor, but KVM can produce better code if it knows the endianness of your machine.

For example, using the Gnu C compiler or the Solaris C compiler, you would write:

```
typedef long long long64;
typedef unsigned long long ulong64;
```

Using Microsoft Visual C, you would write:

```
typedef __int64 long64;
typedef unsigned __int64 ulong64;
```

Your compiler does not support 64-bit integers¹: You must set the preprocessor constant `COMPILER_SUPPORTS_LONG` to zero. You must define exactly one of `BIG_ENDIAN` or `LITTLE_ENDIAN`² to have a non-zero value.

The types `long64` and `ulong64` are defined to be a structure consisting of two fields, each an unsigned long word, named `high` and `low`. The `high` field is first if your machine is big endian; the `low` field is first if your machine is little endian.

You must define the functions shown in Table 9-2. If your platform supports floating point, you must also define the functions shown in Table 9-3.

Any of these functions can be implemented as a macro instead.

TABLE 9-2 Implementing longs

Function or Constant	Java equivalent
<code>long64 ll_mul(long64 a, long64 b);</code>	<code>a * b</code>
<code>long64 ll_div(long64 a, long64 b);</code>	<code>a / b</code>
<code>long64 ll_rem(long64 a, long64 b);</code>	<code>a % b</code>
<code>long64 ll_shl(long64 a, int b);</code>	<code>a << b</code>
<code>long64 ll_shr(long64 a, int b);</code>	<code>a >> b</code>
<code>long64 ll_ushr(long64 a, int b);</code>	<code>a >>> b</code>

TABLE 9-3 Implementing both longs and floats

Function or Constant	Java equivalent
<code>long64 float2ll(float f);</code>	<code>(long)f</code>
<code>long64 double2ll(double d);</code>	<code>(long)d</code>
<code>float ll2float(long64 a);</code>	<code>(float)a</code>
<code>double ll2double(long64 a);</code>	<code>(double)a</code>

1. Or your code must be strictly ANSI standard.

2. See Jonathan Swift, *Gulliver's Travels, Part I: A Voyage to Lilliput*, for more information on the big-endian, little-endian controversy.

9.2 Alignment issues

When an object of Java type `long` or `double` is on the Java stack or in the constant pool, its address will be a multiple of 4.

Some hardware platforms (such as SPARC) require that 64-bit types be aligned so that their address is a multiple of 8.

If your platform requires that 64-bit integers be aligned on 8-byte boundaries, set

```
#define NEED_LONG_ALIGNMENT 1
```

If your platform requires double-precision floating point numbers be aligned on 8-byte boundaries, set

```
#define NEED_DOUBLE_ALIGNMENT 1
```

The compiler can generate better code when these values are 0.

Note – The CLDC standard (CLDC Specification version 1.0) does not support floating point operations. Therefore, all the floating point operations in KVM are turned off by default. The macro `NEED_DOUBLE_ALIGNMENT` needs to be set only in those ports that use floating point operations.

Native Code

KVM does not support the Java Native Interface (JNI). Rather, the native code to be called from the virtual machine must be linked directly into the virtual machine, and must be called using the mechanisms described in this section.

Information for writing your own native functions for KVM is provided in Section 10.2, “Implementing native methods.”

10.1 Native code lookup tables

As part of the build process, you must create the lookup tables that map methods to the corresponding native implementation.

The `JavaCodeCompact` generates these tables automatically. You should use this utility to generate the lookup tables whether or not you are using the other features of `JavaCodeCompact`.

`JavaCodeCompact` is more fully described in Chapter 13. The specific details for creating the file containing the lookup tables can be found in §13.5.

The name of the C function that implements a native method must be the same name that JNI¹ would assign to the native method.

1. See *The Java Native Interface: Programmer's Guide and Specification (Java Series)* by Sheng Liang (Addison Wesley, 1999), for complete information on the JNI naming scheme. This information is available online at <http://java.sun.com/docs/books/jni/index.html>.

10.2 Implementing native methods

WARNING: You should not write native methods unless you have thoroughly read through the implementation and understand its structures. Most of the material in this porting guide is moderately straightforward. The material in this subsection is not!

To avoid the footprint and performance overhead imposed by Java Native Interface (JNI), the KVM reference implementation does not use JNI for native method calls. Native methods must be written extremely carefully. Inattention to detail will cause fatal errors in the virtual machine.

10.2.1 Include files

Your code containing native functions should begin with the line

```
#include <global.h>
```

which causes all include files that are part of KVM to be included. You might also need to #include additional files.

10.2.2 Accessing arguments from native methods

When a native method is called, its arguments are on top of the Java stack. A static method's arguments should be popped from the stack in the *reverse order* from which they were pushed. Figure 10-1 shows an example of this coding style:

```
Java code:
static native void
drawRectangle(int x, int y, int width, int height);

Native implementation:
static void Java_com_sun_kjava_Graphics_drawRectangle() {
    int height = popStack();
    int width = popStack();
    int y = popStack();
    int x = popStack();
    windowSystemDrawRectangle(x, y, width, height);
}
```

FIGURE 10-1 A native method

An instance method (non-static method) must pop the `this` argument off the stack after it has popped the rest of the arguments. *Failing to pop the `this` argument in a native instance method will almost surely cause a fatal error in the virtual machine.*

Table 10-1 shows the macros that should be used to pop arguments off the stack:

TABLE 10-1 Macros for popping arguments from the stack

C type	Macro for popping
<code>char</code> , <code>byte</code> , <code>int</code> , <code>long</code>	<code>popStack()</code>
<code>float</code>	<code>popStackAsType(float)</code>
<code>long64</code> , <code>ulong64</code>	<code>popLong()</code>
<code>double</code>	<code>popDouble()</code>
<code>pointerType</code>	<code>popStackAsType(pointerType)</code>

10.2.3 Returning a result from a native function

If a native method returns a result, it must push that result onto the stack. The native code should use the appropriate macro shown in Table 10-2 to push the result back onto the stack:

TABLE 10-2 Macros for pushing arguments onto the stack

C type	Macro for pushing
<code>char</code> , <code>byte</code> , <code>int</code> , <code>long</code>	<code>pushStack()</code>
<code>float</code>	<code>pushStackAsType(float)</code>
<code>long64</code> , <code>ulong64</code>	<code>pushLong()</code>
<code>double</code>	<code>pushDouble()</code>
<code>pointerType</code>	<code>pushStackAsType(pointerType)</code>

10.2.4 Shortcuts

Some native code uses the macro `topStack` instead of popping the last argument off the stack. It then sets `topStack` to the value it wants to return.

This practice is not encouraged. It should only be used for “one-liners” that access the argument and return the value in a single statement. `pushStack` and `popStack` cannot be used in this case, since C would not guarantee their order of evaluation.

In general, it is safer to pop the value, perform the calculation, and push the value back onto the stack as three separate steps.

10.2.5 Callbacks

Native code cannot call back into Java. KVM provides a mechanism by which native code can alter the interpreter state to begin executing a new piece of code. Upon finishing executing that code, the mechanism can indicate a new C function which should be called.

10.2.6 Exception handling in native code

If the native code needs to throw an error or exception, it should call the function `raiseException(string)` where the `string` argument contains the fully-qualified name (with `'/'` as the package separator) of the exception class or error class.

10.2.7 Useful functions in native code

Other useful functions that a native method might need to call are the following:

- `void fatalError(string)`
The code calls this method to indicate that a serious error has occurred. The `string` argument is a brief explanation of the problem. This method does not return.
- `CLASS getClass(const char *name)`
This method returns the class whose name is the indicated argument. You might want to coerce the return result to be an `INSTANCE_CLASS` or an `ARRAY_CLASS`.
- `INSTANCE instantiateString(const char* string)`
This method converts the given C string into a Java String.
- `char *getStringContents(Instance string)`
The instance argument must be a Java string. It is converted into a null-terminated C string, and returned as the result.
The string is placed into a global buffer. If your code must hold onto this string for any length of time, you must copy the buffer into stack-allocated storage, or allocate space from the Java heap.
- `INSTANCE instantiate(CLASS class)`
Creates a new Java instance of the specified class.
- `ARRAY instantiateArray(ARRAY_CLASS arrayType, long length)`
Creates a Java array of the specified type and length.
- `ARRAY createCharArray(const char* string)`
Creates a Java character array from the C string passed as an argument.

- `char* mallocBytes(long sizeInBytes)`
Allocates a memory block in the garbage-collected heap that is big enough to hold `sizeInBytes` number of bytes. You can create a temporary root (Section 10.2.8, “Garbage collection issues”) to prevent the memory block from being garbage-collected.

10.2.8 Garbage collection issues

The C stack is not scanned when the KVM performs a garbage collection. If your native code allocates new Java objects, you must take special precautions to prevent your new Java objects from being garbage collected inadvertently.

Since the release 1.0.2, KVM includes a compacting garbage collector. Any time that your native code performs an allocation, objects in the Java heap can move. This includes any arguments passed to your native function and any previous heap allocations performed by your native code.

Note – We strongly recommend that you do not write native methods that perform allocation from the Java heap. You greatly increase the chances that your code will have hard-to-find and hard-to-reproduce bugs.

If, for example, you need to create a structure, it is better to create that structure in Java code, and pass it as an argument to the native code.

If your code must perform allocation, it is important that you

- Pop all arguments off the stack before you perform any allocation.
- Push the return value (if any) onto the stack after you have performed any allocation.

The garbage collector can get erroneous results if an allocation occurs while an argument or return value is on the Java stack. The rest of this chapter describes how your code can interact correctly with the garbage collector.

10.2.8.1 Heap Space and Permanent Space

In order to simplify the garbage collector, the KVM’s memory is divided into two spaces: “permanent space” and “heap space”.

All objects created in permanent space are, well, permanent. These objects are

- never freed by the garbage collector,
- never scanned by the garbage collector to see if they contain pointers to other objects,

- never relocated.

Among the objects that are allocated in permanent space are

- class structures,
- Java byte codes,
- method tables,
- field tables,
- interned instances of `java.lang.String` (but not all strings).

These objects never move, and are never freed after they are created.

Structures that have a possibly limited lifetime are allocated in heap space. Among these are

- all Java instances (except for interned `Strings`),
- threads,
- stack chunks.

These structures are liable to move any time an allocation occurs. Your code must following the rules specified in the following subsections to ensure that your code lives happily with the garbage collector.

10.2.8.2 Asserting no allocation

The KVM provides the two macros `ASSERTING_NO_ALLOCATION` and `END_ASSERTING_NO_ALLOCATION` which are used as shown in Figure 10-2 below.

```
ASSERTING_NO_ALLOCATION
    non allocating code
END_ASSERTING_NO_ALLOCATION;
```

FIGURE 10-2 Forbidding garbage collection

If your code is compiled with `INCLUDEDEBUGCODE` set to a non-zero value, then any allocation inside the specified code causes a fatal error.

If you use the macros, make sure that the non-allocating code inside the does not perform a `return`. The macro `END_ASSERTING_NO_ALLOCATION` contains cleanup code that must be executed.

You are encouraged to use these macros to indicate safe regions of code in which heap-allocated objects will not move.

10.2.8.3 Handles

To deal with the fact that heap-allocated objects in the KVM can move, the garbage collector makes use of temporary “handles.” A handle is an indirect pointer to an object. Rather than being the address of the object itself, a handle is the address of a memory location that contains the address of the object.

The memory location that contains the address of the object must not itself be in the Java heap. In general, it is the address of a variable (for global roots) or the address of a location on the C stack (for temporary roots).

- If the object is possibly in the Java heap, then the memory location that contains the address of the object must be registered with the garbage collector. It can either be a temporary root (see §10.2.8.4) or a permanent root (see §10.2.8.5).
- If the object is not in the Java heap, then the handle does not need to be registered with the garbage collector.

All type names in the KVM that end with `_HANDLE` indicate handles. If an argument has a handle as one of its arguments, the argument must be an indirect pointer, and must be registered with the garbage collector if the object could be in the Java heap.

Figure 10-3 below shows an example.

```
CLASS getClassX(CHAR_HANDLE name, int start, int length);

;; Case 1, We are calling it with an argument that is known
;; not to be in the heap.
const char *x = "java/lang/Object";
result = getClassX(&x, 0, strlen(x));

;; Case 2. We are calling it with a heap argument
START_TEMPORARY_ROOTS
    DECLARE_TEMPORARY_ROOT(char *, x, mallocBytes(100));
    sprintf(x, "java/lang/%s", arg);
    result = getClassX(&x, 0, strlen(x));
END_TEMPORARY_ROOTS
```

FIGURE 10-3 Creating a handle

10.2.8.4 Temporary Roots

The most common method is to use `START_TEMPORARY_ROOTS` and `END_TEMPORARY_ROOTS` to delimit a region of code. Within this region of code, the macro

```
DECLARE_TEMPORARY_ROOT(type, variable, value)
```

creates a local variable of the specified type with the specified initial value. The value must either be a pointer to an object in the heap, or it must be a value that is clearly not in the heap (such as `NULL`, a pointer to permanent space, or the like).¹ The value `&variable` is registered with the garbage collector as a temporary root.

You are allowed to change the value of `variable`, provided that any new value is always either a pointer to an object in the heap, or a value that is clearly not in the heap.

The garbage collector ensures whenever a garbage collection occurs, the value of the variable is updated if the value has moved. In addition, `&variable` is a handle, and can be passed as an argument to any function that expects a handle.

Your code must not return. The `END_TEMPORARY_ROOT` contains cleanup code that must be executed.

Figure 10-4 below shows some sample code for a native method that takes a `String` and two integers as arguments, and which must allocate a temporary buffer. If the

```
START_TEMPORARY_ROOTS
    int y = popStack();
    int x = popStack();
    DECLARE_TEMPORARY_ROOT(String_INSTANCE, string,
                           popStackAsType(String_INSTANCE));
    DECLARE_TEMPORARY_ROOT(char*, buffer, mallocBytes(100));
    ;; code that might perform allocation.
END_TEMPORARY_ROOT
```

FIGURE 10-4 Temporary roots

code clearly cannot perform any allocation, then you could instead have written

```
char* buffer = mallocBytes(100);
```

Less commonly used is the macro

```
DECLARE_TEMPORARY_ROOT_FROM_BASE(type, var, value, base)
```

In this case *base* must be a pointer to an object in the heap, and *value* must be a pointer into the middle of the object. The variable *var* is assigned the value *value*. The garbage collector will treat *base* as a root. If *base* is moved by the garbage collector, the value of *var* will be adjusted appropriately.

1. The main purpose of this limitation is that the variable should not have a random integer as its value, and that the variable must be initialized.

10.2.8.5 Global roots

If your code initializes a C variable to point to an object in the Java heap, you can use the code shown in Figure 10-5. There is currently no function for removing a variable from the set of global roots.

```
variable = <value>
MakeGlobalRoot(&(cell **)variable);
```

FIGURE 10-5 Creating a global root

This code ensures that the garbage collector knows that the specified variable contains a value that must be protected from garbage collection. If the garbage collector moves the object, the variable is updated to point to the new value.

10.2.8.6 Debugging your native code

A special garbage collector is provided to help you debug your native code and to ensure that it does not have any garbage collection problems. You access this garbage collector by replacing the file `collector.c` with `collectorDebug.c`. In addition, you should set the compiler flags `INCLUDEDEBUGCODE` and `EXCESSIVE_GARBAGE_COLLECTION` to 1.

This replaces the compact-in-place garbage collector with a 10-space Cheney style¹ garbage collection algorithm. A garbage-collection will occur on every allocation, and also on some operations that might have allocated but didn't. Every object moves on every garbage collection. In addition, this code makes use of memory-protection so that any attempts to read or write a bad pointer will generate a memory fault.

This code makes use of the following implementation-dependent functions:

```
void* allocateVirtualMemory_md(long size);
void freeVirtualMemory_md(void *address, long size);
void protectVirtualMemory_md(void *address, long size,
                             int protection);
```

Implementations of these three functions for Windows and for Unix are provided.

1. C.J. Cheney. A non-recursive list compacting algorithm. *Communications of the ACM*, 13(11):677-8, November 1970.

10.2.8.7 Two-space Cheney garbage collector

The file `collectorDebug.c` (see §10.2.8.6) also includes an implementation of a two-space non-debugging Cheney garbage collector. You get this implementation by setting the compiler flag `CHENEY_TWO_SPACE` to a non-zero value.

The Cheney collector is smaller and faster than the standard garbage collector. However, it uses twice as much heap space. If your implementations has a lot of available memory, but needs a faster garbage collector, you might consider using this garbage collector.

This collector is not supported by Sun, and is provided as is.

10.2.9 Initialization and reinitialization of global variables

Generally, the C language guarantees that all global and static variables are initialized to 0 (zero).

The current implementation is designed to work within an embedded environment. For example, on the PalmOS, the user can start the virtual machine, exit a program, and then restart the virtual machine with a different set of arguments. There is no re-initialization of global or static variables between the two runs.

In general, your code cannot assume the initial value of any variable. You have several options for determining when it is necessary to perform one-time only initialization.

- You can use the function `initializeNativeMethods()` to either initialize your variables, or to set a flag indicating that initialization needs to be performed.
- If a private native method is called as part of static initialization of a class, the method's native implementation will be called the first time the class is used. The native implementation can perform any initialization necessary for the class.
- If a variable is part of the global root set (see `MakeGlobalRoot()` above), its value is guaranteed to be 0 the next time that the virtual machine is run.

10.3 Asynchronous native methods

From the operating system viewpoint, KVM is just one process (C program) with only one native thread of execution. The multithreading capabilities of KVM have been implemented entirely in software without utilizing the possible multitasking capabilities of the underlying operating system. This approach not only makes the

virtual machine highly portable and independent of the operating system, but also greatly simplifies the virtual machine design and improves the readability of the codebase, as the virtual machine designer does not have to worry about mutual exclusion and other problems typically associated with multithreaded software.

However, an unfortunate side effect of the approach described above is that by default, all native methods in KVM are “blocking.” This means that when a native function is called from the virtual machine, all the threads in the VM stop executing until the native method completes execution.

As a general guideline, all the native functions called from KVM should be written so that they complete their execution as soon as possible. However, in many environments this is not desirable or fully possible. For this reason, KVM includes an implementation of “asynchronous native methods” described below.

10.3.1 Design of asynchronous methods

The standard implementation of KVM runs as a single “task” from the operating system’s point of view. If a native method performs an operation that can block, the entire KVM blocks.

Asynchronous native methods are intended to solve this problem. When such a native method is called, the operation is performed “off-line” in an implementation-dependent manner. Other Java threads can continue running normally. When the native call finishes, the Java thread that originally called the native method continues.

To use asynchronous native methods, you must include

```
#define ASYNCHRONOUS_NATIVE_METHODS 1
```

in your machine-dependent include file.

Asynchronous native methods cannot be defined in the same file as normal native methods. In addition to their normal includes, they must also add the include the file `async.h`.

Asynchronous methods should always have the following form:

```
ASYNC_FUNCTION_START(functionname)  
    code  
ASYNC_FUNCTION_END
```

Your code must never use `pushStack()`, `popStack()`, `topStack`, or any macro or function that references the stack pointer, the frame pointer, or the current thread. Instead, you must use the alternative macros shown in Table 10-3.

TABLE 10-3 Macros used in asynchronous methods

Native function macro	Asynchronous native function macro
<code>popStack</code>	<code>ASYNC_popStack</code>
<code>pushStack</code>	<code>ASYNC_pushStack</code>
<code>popLong</code>	<code>ASYNC_popLong</code>
<code>pushLong</code>	<code>ASYNC_pushLong</code>
<code>popStackAsType</code>	<code>ASYNC_popStackAsType</code>
<code>pushStackAsType</code>	<code>ASYNC_pushStackAsType</code>
<code>raiseException</code>	<code>ASYNC_raiseException</code>
<code>topStack</code>	do not use this macro

In addition, your code must not perform a “return.” It must complete through the end, since `ASYNC_FUNCTION_END` may generate some necessary cleanup code.

All the macros in Table 10-3 have been designed so that if the symbol `ASYNCHRONOUS_NATIVE_METHODS` is 0, the asynchronous method compiles into a normal native method.

It is also important to note that unlike regular native methods, asynchronous native methods *cannot allocate any memory from the Java heap*. Because of this limitation, extra caution is often necessary when writing asynchronous native methods, since many internal routines in KVM may indirectly allocate memory from the Java heap.

Note – We repeat that asynchronous native methods must not allocate memory from the Java heap. Make sure that you read the paragraph above.

If you use asynchronous native methods, you must define the following machine-specific functions.

- `void yield_md()`
Pause this operating system task momentarily and allow other tasks to run.
- `void CallAsyncNativeFunction_md(ASYNCIOCB *, void(f*)(ASYNCIOCB *))`
Call an asynchronous native function. This function is called by the `ASYNC_FUNCTION_START` macro to start a new asynchronous function. The function takes as a parameter a data structure that is used by the garbage collector

to keep up to date object pointers used by the native code, and a function to call. This function will typically start a new native thread and have that call the supplied function with the ASYNCIOCB as its parameter.

- `enterSystemCriticalSection()`
`exitSystemCriticalSection()`
 Enter or exit a critical section. The operating system must guarantee that at most one operating system task is allowed to be inside the critical section at a time.

10.3.2 Implementation of asynchronous methods

We envision two possible implementations of asynchronous methods.

In the current reference implementation, the function `CallAsyncNativeFunction_md` spawns off a separate operating system task which performs the indicated function. For example, in a Posix implementation one could use `pthread_create`.

Figure 10-6 below shows one possible implementation of a method
`int readBytes(byte[] dst, int offset, int length)`
 using this style of asynchronous native methods.

```
ASYNC_FUNCTION_START(ReadBytes)
    long    length = ASYNC_popStack();
    long    offset = ASYNC_popStack();
    BYTEARRAY dst = ASYNC_popStackAsType(BYTEARRAY)
    INSTANCE instance = ASYHNC_popStackAsType(INSTANCE) /* this*/
    long fd = getFD(instance);
    ASYNC_enableGarbageCollection();
    length = read(fd, dst->bdata + offset, length);
    ASYNC_disableGarbageCollection();
    ASYNC_pushStack((length == 0) ? -1 : length);
ASYNC_FUNCTION_END
```

FIGURE 10-6 Asynchronous implementation of `ReadBytes`

In an alternative implementation, `CallAsyncNativeFunction_md` simply calls the function `f` directly. It assumes that the function `f` starts an operation, but does not wait for its completion. The operating system is required to provide some sort of interrupt or callback to indicate when the operation is complete.

The second implementation is far more operating system-dependent. It might be impossible to write native methods that can work both synchronously and asynchronously, depending on the value of a flag.

Refer to Section 11.1.4, “Asynchronous notification,” for further information on writing asynchronous code.

```

static void ReadBytes(THREAD thisThread)
{
    long    length = ASYNC_popStack();
    long    offset = ASYNC_popStack();
    BYTEARRAY dst = ASYNC_popStackAsType(BYTEARRAY);
    INSTANCE instance = ASYNC_popStackAsType(INSTANCE);
    long fd = getFD(instance);
    THREAD thisThread = CurrentThread;
    /* Call OS to perform I/O. Perform callback when done. */
    AsyncRead(fd, p + offset, length, ReadBytesDone, thisThread);
}

/* Callback function when I/O is finished */
static void ReadBytesDone(void *parm, int length)
{
    THREAD thisThread = (THREAD)parm;
    ASYNC_pushStack((length == 0) ? -1 : length);
    ASYNC_RESUME_THREAD();
}

```

FIGURE 10-7 Alternative asynchronous implementation of ReadBytes

Event Handling

11.1 High-level description

The Java Virtual Machine Specification does not define how the virtual machine interacts with events that arrive from the host operating system or from the target device. The KVM implementation, however, provides a variety of mechanisms that have been designed to facilitate the integration of the KVM with the event system mechanisms of the host operating system or device.

There are four ways in which notification and handling of events can be accomplished in KVM:

1. Synchronous notification (blocking).
2. Polling in Java code.
3. Polling in the bytecode interpreter.
4. Asynchronous notification.

Different solutions may be appropriate for different parts of the KVM, depending on which user interface libraries are supported, what kinds of networking libraries are used, etc.

11.1.1 Synchronous notification (blocking)

By synchronous notification we refer to a situation in which the KVM performs event handling by calling a native I/O or event system function directly from the virtual machine. Since the KVM has only one physical thread of control inside the virtual machine, no other Java threads can be processed while the native function is

being executed, and no VM system functions like garbage collection can occur either. This is the simplest form of event notification, but there are many situations in which this solution is quite acceptable, provided that the person designing the native functions is careful enough to keep the native functions as short and efficient as possible.

For instance, writing a datagram into the network can typically be performed efficiently using this approach, since typically the datagram is sent to a network stack that contains a buffer and the time spent waiting for the event to complete is very small. In contrast, reading a datagram is often a very different story, and is often handled better using the other solutions described below. Using a native function to wait until a whole datagram is received would block the whole KVM while the read operation is in progress.

11.1.2 Polling in Java code

Often event handling can be implemented efficiently using a combination of native and Java code. This is a simple way to allow other Java threads to execute while waiting for an event to complete. When using this approach, a polling Java loop is normally put somewhere in the Java runtime libraries so that the loop is hidden from applications. The normal procedure is for the runtime library to initiate a short native I/O operation and then repeatedly query the status of the I/O operation until it is finished. The polling Java code loop should always contain a call to `Thread.yield` so that other Java threads can be allowed to run efficiently.

This method of waiting for event notification is very easy to implement and is free of any complexities typically associated with genuinely asynchronous threads (such as requiring critical sections, semaphores or monitors.) There are two disadvantages with this design. First, CPU cycles are needed to perform the Java-level polling that could otherwise be used to run application code (although the overhead is usually very small.) Second, due to the interpretation overhead, there may be some extra latency associated with event notification (especially if you forget to call `Thread.yield` in the polling Java code loop.) Again, this overhead is usually negligible in all but most time-critical applications.

11.1.3 Polling in the bytecode interpreter

The third approach to implement event handling is to use the bytecode interpreter periodically make calls to the native event handling operations. This approach is a variation of the synchronous notification approach described above. This approach was originally used extensively in the KVM, e.g., to implement GUI event handling for the Palm platform.

In this approach, a native event handling function is called periodically from the interpreter loop. For performance reasons this is not normally done before every bytecode, but every few hundred bytecodes or so. This way the cost of performing event handling is well amortized. By changing the number of bytecodes executed before calling the event handling code, the virtual machine designer can control the latency of event delivery versus the CPU time spent looking for a new event. The smaller the number, the smaller latency and the larger CPU overhead. A large number reduces CPU overhead but increases the latency in event handling.

The advantage of this approach is that the cost in performance is less than polling in Java, and the event notification latency is more predictable and controllable. The way this approach works is closely related to asynchronous notification described in the next subsection.

11.1.4 Asynchronous notification

The original KVM implementation supported only the three event handling implementations discussed above. However, in order to support truly asynchronous event handling, some new mechanisms have been introduced.

By asynchronous notification we refer to a situation in which event handling can occur in parallel while the virtual machine continues its execution. This is generally the most efficient event handling approach and will typically result in a very low notification latency. However, this approach generally requires that the underlying operating system provides the appropriate facilities for implementing asynchronous event handling. Such facilities may not be available in all operating systems. Also, this approach is quite a bit more complex to implement, as the virtual machine designer must be aware of possible locking and mutual exclusion issues. The reference implementation provides some examples that can be used as a starting point when implementing more device-specific event handling operations.

The general procedure in asynchronous notification is as follows. A thread calls a native function to start an I/O operation. The native code then suspends the thread's execution and immediately exits back to the interpreter loop, letting other threads continue execution. The interpreter then selects a new thread to run. Some time later an asynchronous event occurs and as a result some native code is executed which resumes the suspended thread. The interpreter then restarts the execution of the thread that had been waiting for an event to occur.

At the implementation level, there are two ways to implement such asynchronous notification. One is to use native (operating system) threads, and the other is to use some kind of software interrupt, callback routine or a polling routine.

In the first case, before the native function is called and the Java thread is suspended, a new operating system thread is created (or reawakened) and it is this thread which enters the native function. There is now an additional native thread of

control running inside the virtual machine. After the native I/O thread is started, the order of execution inside the virtual machine is no longer fully deterministic, but depends on the occurrence of external events. Typically, the original thread starts executing another Java thread in the interpreter loop, and the new thread starts the I/O operation with what is almost always a blocking I/O operation to the operating system.

It is important to note that the native I/O function will execute out of context meaning that the context of the virtual machine will be a different thread. A special set of C macros have been written that will hide this fact for the most part, but special care should be taken to be sure that no contextual pointers are used in this routine. When the blocking call is finished the native I/O thread resumes execution and unblocks the Java thread it was representing. The Java thread is then rescheduled, and the native I/O thread is either destroyed, or placed in a dormant state until it needs to be used again. The Win32 port of the KVM reference implementation does this by creating a pool of I/O threads that are reused when I/O is to be performed.

The second implementation of asynchronous event handling can be done by utilizing callback functions associated with I/O requests. Here the native code is entered using the normal interpreter thread, I/O is started and then when the I/O operation is completed a callback routine is called by the operating system and the Java thread is unsuspended. In this scenario the native code is split into two routines, the first being a routine that starts the I/O operation and the second invoked when I/O is completed. In this case the first routine runs in the context of the calling Java thread, and the second one does not.

The final, less efficient variation of asynchronous event handling is where the I/O routine is polled for completion by the interpreter loop. This is very similar to the callback approach except that the second routine is called repeatedly by the interpreter to check if the I/O has finished. Eventually when the I/O operation has completed the routine unblocks the waiting Java thread. This calling of the native code by the interpreter is always done even when there are no pending events, and the native code must determine what Java threads should be restarted.

Synchronization issues. It is very important to remember that in the cases where a separate native event handling thread or callback routine is used, the code for event handling may interrupt the virtual machine at any point. Therefore, the person porting the virtual machine must remember to add critical sections, monitors or semaphores to all locations where the program may be manipulating common data structures and a possible mutual exclusion problem might occur. The most obvious shared data structures are the queues of suspended and active Java threads. These are always manipulated using special routine in the virtual machine that is already properly synchronized. If there are any other shared data structures they must be synchronized in the native code. Failure to do this correctly will produce spurious bugs that are very hard to debug.

11.2 Parameter passing and garbage collection issues

When native event handling code is called, its parameters will be on the stack for the calling Java thread. These are popped off the stack by the native code, and the if there is a result value to be returned this is pushed onto the Java stack just prior to resuming the execution of the thread. Native parameter passing issues have been discussed in Chapter 10.

Because native event handling code can access object memory, there are possible garbage collection issues especially when running long, asynchronous I/O operations. In general, the garbage collector is prevented from running when there is any native code is running. This is a problem when certain long I/O operations are performed. The most obvious case is waiting for a incoming network request. To solve this problem two functions called `decrementAsyncCount` and `incrementAsyncCount` are provided. The first allows the garbage collector to start, and the second prevents the collector from starting, and waits for it to stop if it was running.

It should be noted that if an object reference is passed to a native method, but no other reference to it exists in Java code after the call to `incrementAsyncCount`, the object could be reclaimed accidentally by the garbage collector. It is hard to think of a realistic scenario where this could occur, but the possibility should be kept in mind. A possible example of such code is the following:

```
native read(byte[]);
void skipBytes(int n) {
    read(new byte[n]);
}
```

Here the only reference to the byte array object exists on the parameter stack to the native function. If the native code calls `incrementAsyncCount` after popping the parameter from the stack the array could be garbage collected.

11.3 Implementation in KVM

The event handling implementation in KVM is composed of two main layers that both need to be taken into account when porting the KVM onto new hardware platforms.

At the top of the interpreter loop is the following code (starting from KVM 1.0.2, this code is actually located in macros):

```
if (isTimeToReschedule())  
    reschedule();
```

The standard rescheduling code performs the following operations.

1. Checks to see if there are any active Java threads and stops the VM if there are none.
2. Checks to see if enough time has passed to allow a thread that was waiting for a specific time to be restarted. If there is such a thread, it is automatically restarted.
3. Checks to see if any I/O events have occurred and where appropriate it allows the relevant threads to contend for CPU time
4. Attempts to switch to another thread.

For performance reasons, the operations above are implemented as macros that are, by default, defined in `VmCommon/h/events.h`. It is here that device-specific event handling code can be placed. By default, the `isTimeToReschedule` macro decrements a global counter and tests for it being zero. When it is zero the second macro is executed. The idea here is for the `reschedule` to be executed only once for a fairly large number of bytecode executions. As the name implies, `reschedule` is where the thread context switching is done, if necessary.

The second layer in event handling implementation is the function

```
GetAndStoreNextKVMEvent(bool_t forever, ulong64 waitUntil)
```

If a new event is available from the host operating system, this function must call a special function called `StoreKVMEvent` to make the details of the event available to the KVM. If no new events are available from the host operating system, then the function can simply return.

The arguments to the `GetAndStoreNextKVMEvent` function are as follows:

- If the `forever` argument is `TRUE`, this function should wait for as long as necessary for an event to occur (used for battery conservation as described below.)
- If the `forever` argument is `FALSE`, this function should wait until at most `waitUntil` for an event to occur.

Some battery conservation features have been included in the reference implementation of these functions. This is to pass to the event checking function the “forever” flag or the maximum wait time. If there are no pending events, the native implementation of the `GetAndStoreNextKVMEvent` function can then put the device “to sleep” until the next event occurs. Battery conservation issues have been discussed in more detail in the next subsection.

11.4 Battery power conservation

Most KVM target devices are battery-operated, and the manufacturers of these devices are typically extremely concerned of excessive battery power consumption. To minimize battery usage, KVM is designed to stop the KVM interpreter loop from running whenever there are no active Java threads in the virtual machine and when the virtual machine is waiting for external events to occur. This requires support from the underlying operating system, however.

In order to take advantage of the power conservation features, you must port the following low-level event reading function

```
GetAndStoreNextKVMEvent(bool_t forever, ulong64 waitUntil)
```

so that it calls the host system specific sleep/hibernation features when the virtual machine calls this function with the `forever` argument set `TRUE`. The KVM has been designed to automatically call this function with the `forever` argument set `TRUE` if the virtual machine has nothing else to do at the time.

This allows the native implementation of the event reading function to call the appropriate device-specific sleep/hibernation features until the next native event occurs.

Additionally, the macro `SLEEP_UNTIL(wakeupTime)` should be defined in such a fashion that the target device goes to sleep until `wakeupTime` milliseconds has passed.

Class File Verification

12.1 Overview

The class file verifier supported by Java 2 Standard Edition (J2SE) is not suitable for small, resource-constrained devices. The J2SE verifier requires a minimum of 50 kB binary code space, and at least 30-100 kB of dynamic RAM at run time. In addition, the CPU power needed to perform the iterative dataflow algorithm in the standard JDK verifier can be substantial.

We have designed and implemented a new, two-phase class file verifier that is significantly smaller than the existing J2SE verifier. The runtime part of the new verifier requires about 15 kB of Intel x86 binary code and only a few hundred bytes of dynamic RAM at run time for typical class files. The runtime verifier performs a linear scan of the byte code, without the need of a costly iterative dataflow algorithm. The new verifier is especially suitable for KVM, a small-footprint Java virtual machine for resource-constrained devices.

The new class file verifier operates in two phases, as illustrated in Figure 12-1:

- First, Java class files have to be run through a special *preverifier* tool in order to augment the class files with additional attributes to speed up runtime verification. The preverification phase is typically performed on a development workstation, where the application developer writes and compiles the applications.
- At runtime, the runtime verifier component in the KVM utilizes the additional attributes generated by the preverifier to perform the actual class file verification efficiently.

Development workstation

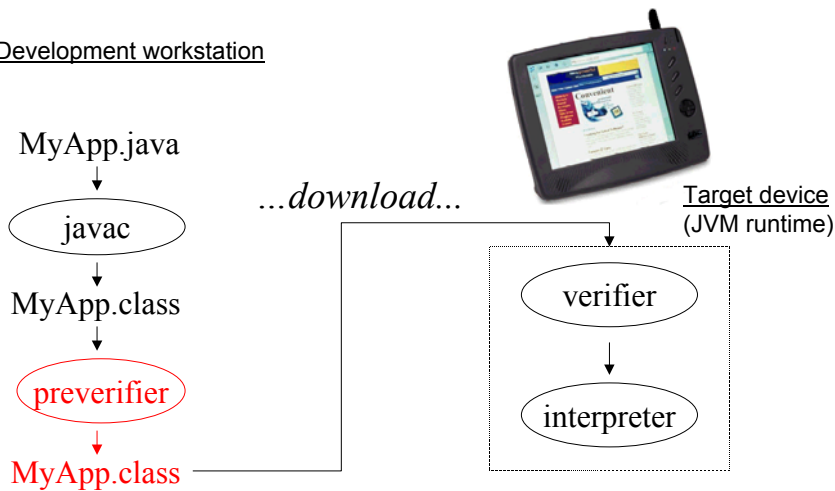


FIGURE 12-1 Two-phase verification

The runtime class file verifier requires all the subroutines to be inlined, so that class files contain no `jsr`, `jsr _w`, `ret`, or `wide ret` instructions. Additionally, the runtime verifier requires the methods in class files to contain special `StackMap` attributes. The preverifier tool performs these modifications to normal class files generated by a Java compiler such as `javac`. A transformed class file is still a valid J2SE class file, but with additional attributes that allow verification to be carried out efficiently at run time.

Note – In the future, `javac` (the Java compiler) may be modified to perform these changes automatically. In that case, the preverifier tool will no longer be necessary.

The preverifier tool shipped with the KVM release is a C program that contains code extracted from the JDK 1.1.8 virtual machine implementation as well as code specifically written for inlining subroutines and inserting the `StackMap` attributes. The program compiles and runs on Windows, Solaris and Linux, and can be ported to other development platforms relatively easily.

12.2 Using the preverifier

The preverification phase is usually performed at application development time on a development workstation. The preverifier is used as follows. If, for example, you normally compile `Foo.java` using `javac` like this:

```
javac -classpath kvm/classes Foo.java
```

When using the preverifier, you typically place the output of `javac` in a separate directory and then transform the resulting class files using the preverifier:

```
javac -classpath kvm/classes -d mydir Foo.java
preverify -classpath kvm/classes -d . mydir
```

The above preverifier command transforms all class files under `mydir/` and places the transformed class files in the current directory (as specified by the `-d` option).

Makefiles in the KVM distribution invoke the preverifier automatically.

12.2.1 General form

More generally, the preverifier is invoked as follows:

```
preverify <options> <input files>
```

Preverifier options and accepted input file formats are explained in more detail below.

12.2.2 Preverifier options

The preverifier accepts a number of arguments and options.

`-classpath <directories> | <JAR files>`

- Directories or JAR file(s) in which the KVM/CLDC Java library classes are located. The directory separator is platform-specific. On Solaris a colon is used. On Win32 a semicolon is used. The JAR file specified must be in a valid Java Archive format and must end with either `“.jar”`, `“.JAR”`, `“.zip”` or `“.ZIP”` suffix.

`-d <directory>`

- The directory in which output classes will be written. The default output directory is `./output`.

`-cldc`

- If specified, this option checks for the usage of those Java language features that are prohibited by the CLDC Specification, including floating point operations, finalizers, and the use of native methods in application classes. An error is reported if any of these features are detected in any of the input files.

`-nofinalize`

- This option checks for the use of finalizers in application classes. When this option is specified, an error is reported if finalizers are detected in any of the input files.

`-nonative`

- This option checks for the use of native methods in application classes. When this option is specified, an error is reported if native methods are detected in any of the input files.

`-nofp`

- This option checks for the use of floating point operations in application classes. When this option is specified, an error is reported if floating point operations are detected in any of the input files.

`@<filename>`

- The name of a text file from which command line arguments will be read.

Note – When the command line arguments are read from a file, parameters must all be specified on a single line and the parameters to the “`-classpath`” and “`-d`” options must be enclosed within double quotes. When the corresponding options are used from the command line, quotes are not required (unless the directory/file name parameter contains spaces.)

For example, the contents of `<filename>` under Win32 may appear as follows:

```
-classpath "api/classes; aaa bbb ccc/samples/classes" -d "output"
-verbose HelloWorld1 HelloWorld2 HelloWorld3
```

12.2.3 Supported input file formats

The preverifier can accept input files in three different formats:

- individual Java class files
- directories containing Java class files
- JAR files containing Java class files.

To preverify a single class file or a number of class files, simply include the class file(s) after the command line options:

```
preverify -classpath kvm/classes File1 File2 ...
```

To preverify all the Java class files contained in a directory or set of directories, invoke the preverifier tool as follows:

```
preverify -classpath kvm/classes dir1 dir2 ...
```

To preverify all the Java class files contained in one or more JAR files, invoke the preverifier tool as follows:

```
preverify -classpath kvm/classes Jar1.jar Jar2.jar ...
```

Any combination of individual class files, directories or JAR files should be possible.

Obviously, the library classes can also be contained in a JAR/ZIP file, as illustrated by the line below:

```
preverify -classpath classes.zip File1 File2 ...
```

Output is generated differently depending on input parameters. If individual files are specified, the preverifier tool performs preverification separately for each input file. For each directory name, the preverifier recursively transforms every class file under that directory. The JAR file handling is discussed in the next section.

Note – A non-zero error status is returned if preverification fails for any reason.

12.2.4 JAR support in preverifier (since KVM 1.0.2)

Since KVM 1.0.2, the preverifier tool provided with the KVM allows input files to be provided as a JAR file. Given a JAR file that contains un-preverified Java class files, the preverifier tool will automatically generate an identical JAR file containing preverified class files.

This is performed as follows: First, the preverifier will check the file extension (".jar", ".JAR", ".zip" or ".ZIP" file suffixes are acceptable) and validate that the file is in valid Java Archive format. Then, the class files will be extracted from the JAR file. For each class name extracted from the JAR file, the preverifier tool will perform the necessary transformations, and will then store the output file into a temporary directory `tmpdir`. After all the class files have been transformed successfully, a new JAR file with the same name will be created under `<output>` directory containing all the verified classes previously stored in `tmpdir`.

If the preverifier is run in non-verbose mode, any errors that may have occurred during the JAR creation will be logged in the `<output>/jarlog.txt` file, where `<output>` refers to the directory in which output classes will be written. If no errors occur during JAR creation, the `<output>/jarlog.txt` file will be removed. Directory `tmpdir` is also removed after the JAR file creation.

Note – When preverifying class files contained in JAR files, the preverifier tool will internally call the standard JAR tool to repackage the output files into a new JAR file. To accomplish this, the standard JAR tool must be accessible on your file path.

12.3 Porting the verifier

Runtime part. The runtime part of the new verifier does not generally require any porting efforts, as it is closely integrated with the rest of the virtual machine, and is implemented in portable C code.

Preverifier part. The preverifier is also written in C. By default, the preverifier is available for Windows and Solaris, but it should be relatively easy to compile it to run on other operating systems as well. Note that the preverifier codebase is derived from the “Classic” JVM, so the preverifier implementation looks quite different from the rest of the KVM codebase.

12.3.1 Compiling the preverifier

The sources for the preverifier are in the directory `tools/preverifier/src`.

On Solaris, you can build the preverifier by typing the “`gnumake`” command in the `tools/preverifier/build/solaris` directory. This compiles and links all `.c` files in the `tools/preverifier/src` directory, and places the resulting executable file in the `tools/preverifier/build/solaris` directory.

On Win32, you can build the preverifier by typing the “`gnumake`” command in the `tools/preverifier/build/win32` directory. This compiles and links all `.c` files in the `tools/preverifier/src` subdirectory, and places the resulting executable file in the `tools/preverifier/build/win32` directory.

JavaCodeCompact (JCC)

KVM supports the *JavaCodeCompact* (JCC) utility (also known as the class *prelinker*, *preloader* or *ROMizer*). This utility allows Java classes to be linked directly in the virtual machine, reducing VM startup time considerably.

At the implementation level, the JavaCodeCompact utility combines Java class files and produces a C file that can be compiled and linked with the Java virtual machine.

In conventional class loading, you use `javac` to compile Java source files into Java class files. These class files are loaded into a Java system, either individually, or as part of a jar archive file. Upon demand, the class loading mechanism resolves references to other class definitions.

JavaCodeCompact provides an alternative means of program linking and symbol resolution, one that provides a less-flexible model of program building, but which helps reduce the VM's bandwidth and memory requirements.

JavaCodeCompact can:

- combine multiple input files
- determine an object instance's layout and size
- load only designated class members, discarding others.

13.1 JavaCodeCompact options

JavaCodeCompact accepts a large number of arguments and options. Only the options currently supported by KVM are given below.

- *filename*

Designates the name of a file to be used as input, the contents of which should be included in the output. File names with a `.class` suffix are read as single-class files.

File names with `.jar` or `.zip` suffixes are read as Zip files. Class files contained as elements of these files are read. Other elements are silently ignored.

- `-o output filename`

Designates the name of the output file to be produced. In the absence of this option, a file is produced with the name `ROMjava.c`.

- `-nq`

Prevents `JavaCodeCompact` from converting the byte codes into their “quickenized” form. This option is currently required by KVM.

- `-classpath path`

Specifies the path `JavaCodeCompact` uses to look up classes. Directories and zip files are separated by the delimiting character defined by the Java constant `java.io.File.pathSeparatorChar`. This character is usually a colon on the Unix platform, and a semicolon on the Windows platform.

Multiple classpath options are cumulative, and are searched left-to-right. This option is used in conjunction with the `-c` cumulative-linking option, and with the `-memberlist` selective-linking option.

- `-memberlist filename`

Performs selective loading as directed by the indicated file. This file is an ASCII file, as produced by `JavaFilter`, containing the names of classes and class members.

- `-v`

Turns up the verbosity of the linking process. This option is cumulative. Currently up to three levels of verbosity are understood. This option is only of interest as a debugging aid.

- `-arch Architecture`

Specify the architecture for which you are generating a romized image. If you are using `JavaCodeCompact` for the PalmOS, you must specify `PALM` as the architecture; otherwise, you must specify `KVM` as the architecture.

13.2 Porting JavaCodeCompact

With one exception, `JavaCodeCompact` outputs C code that is completely platform-independent.

To initialize a variable that is `final static long` or `final static double`, `JavaCodeCompact` performs the appropriate initialization using the two macros:

```
ROM_STATIC_LONG(high-32-bits, low-32-bits)
ROM_STATIC_DOUBLE(high-32-bits, low-32-bits)
```

If you have initialized either the compiler `BIG_ENDIAN` or `LITTLE_ENDIAN` to a non-zero value, the file `src/VmCommon/h/rom.h` generates default values for these macros.

If you have not defined `BIG_ENDIAN` or `LITTLE_ENDIAN`, or if for some reason the macros defined in `rom.h` are inappropriate for your platform, you should create appropriate definitions for `ROM_STATIC_LONG` and/or `ROM_STATIC_DOUBLE` in a platform-dependent location.

There are no other known platform or port dependencies.

13.3 Compiling JavaCodeCompact

The sources for JavaCodeCompact are in the directory `tools/jcc/src`.

On Unix and Windows machines, you compile JavaCodeCompact by typing the command “`gnumake`” in the `tools/jcc/` directory. This compiles all `.java` files in the `tools/jcc/src` subdirectory, and places the resulting compiled file in the `tools/jcc/classes` directory.

You may need to make modifications to this file to indicate the location of your `javac` compiler.

13.4 JavaCodeCompact files

The directory `tools/jcc` contains a `Makefile` that shows all the steps necessary to execute JavaCodeCompact. This `Makefile` currently has three targets:

```
unix
windows
palm
```

each of which can be used to create all the files necessary for that platform.

On the `unix` and `windows` platforms, two files are created:

```
ROMjavaPlatform.c
nativeFunctionTablePlatform.c
```

The first file contains the C data structures that correspond to the classes in the zip file. The second file contains tables necessary for using native functions (see §10.1). This second file should be compiled and linked into KVM whether or not you are planning to use the other features of the JavaCodeCompact utility.

On the Palm, several files are created:

```
nativeFunctionTablePalm.c
nativeRelocationPalm.c
kvm/VmPilot/build/bin/PalmROM1001.bin
...
kvm/VmPilot/build/bin/PalmROM1010.bin
```

The file `nativeFunctionTable.c` again contains tables necessary for using native functions. It should be compiled and linked into KVM whether or not you are planning to use the other features of the `JavaCodeCompact` utility. The file `nativeRelocationPalm.c` contains relocation information needed to execute native methods. The directory `kvm/VmPilot/build/bin` contains a set of Palm resource files that must be included in your `kvm.prc` file.

13.5 Executing JavaCodeCompact

The `JavaCodeCompact` utility is used to build the platform-specific file `nativeFunctionTablePlatform.c`, which contains tables necessary for calling native methods.

This file must be built even if you are not using the ability of `JavaCodeCompact` to pre-load classes for you.

If you are not using `JavaCodeCompact`, you may skip Step 4 below.

The simplest method for using the `JavaCodeCompact` utility is to either use the Makefile provided or to modify it for your platform. The following lists the steps that the makefile performs:

1. Compile all the .java files in the `api/src` directory. The resulting class files are verified and merged into a single zip file `classes.zip`. This zip file is copied to the `tools/jcc` directory.
2. Compile the sources for JCC as described in §13.3 above.
3. Copy `classes.zip` to `classesPlatform.zip`. Remove from this platform-dependent zip file any classes or packages that should not be used on your platform.
- 4a.[Not Palm] Execute your system's equivalent of the following command in the `jcc` directory:

```
env CLASSPATH=classes \
  JavaCodeCompact -nq -arch KVM \
  -o ROMjavaPlatform.c classesPlatform.zip
```

The “`env CLASSPATH=classes`” sets an environment variable indicating that

the code for executing `JavaCodeCompact` can be found in the subdirectory called `classes`. Next on the command line is the name of the class whose main method is to be executed (`JavaCodeCompact`), and the arguments to that method.

4b.[Palm] You should instead execute the following two commands:

```
env CLASSPATH=classes \
  JavaCodeCompact -nq -arch Palm \
  -o ROMjavaPalm.c classesPalm.zip
env CLASSPATH=classes \
  JavaCodeCompact -nq -arch Palm \
  -imageAttribute relocating
  -o nativeRelocationPalm.c classesPalm.zip
```

The file `nativeRelocationPalm.c` is included as a source file in your build. The file should be compiled and executed as follows:

```
cc -I../kvm/VmPilot/h -I../ikvmvm/VmCommon/h \
  -DRELOCATABLE_ROM -DROMIZING ROMjavaPalm.c \
  -o ROMjavaPalm
```

The resulting executable `ROMjavaPalm` is executed as follows:

```
ROMjavaPalm ../../kvm/VmPilot/build/bin/PalmROM
```

creates the resource files in the indicated directory.

5. Execute your system's equivalent of the following command in the `jcc` directory:

```
env CLASSPATH=classes \
  JavaCodeCompact -nq -arch KVM_Native
  -o nativeFunctionTablePlatform.c classesPlatform.zip
```

This command creates the file containing the native function tables necessary to link native methods to the corresponding C code.

6. Recompile all the sources for KVM. You must ensure that the preprocessor macro `ROMIZING` is set to a non-zero integer value. You must also ensure that the file `ROMjavaPlatform.c` (non Palm) or `nativeRelocationPalm.c` (Palm) is included as one of your source files.

The resulting kvm image will include, pre-loaded, all of the class files that were in the original `classesPlatform.zip` file.

13.6 Limitations

The current implementation of `JavaCodeCompact` requires that the class files that you compact constitute a “transitive closure.” If class A is compacted, and class A’s constant pool references class B, then class B must also be included as part of the compaction.

Class A includes Class B in its constant pool if any of the following conditions are true:

- Class A is a direct subclass of class B, or class A directly implements class B.
- Class A creates an instance of class B, or an array of class B.
- Class A calls a method that is defined in class B.
- Class A checks to see if an object is an instance of type B, or casts an object to type B.

Note that the following do not cause class B to be included in class A’s constant pool. Under certain circumstances, it may be possible to compact A without also compacting B.

- Class A has an instance variable of type B
- Class A has a method whose argument or return type includes type B in its signature.
- Class A creates an instance of class B using the `Class.forName()` method.

`JavaCodeCompact` will fail and give you an error message if you fail to include a class file that it requires.

Java Application Manager (JAM)

A central requirement for KVM in most target devices is to be able to execute applications that have been downloaded dynamically from the network. Once downloaded, the user commonly wants to use the applications several times before deleting them. The process of downloading, installing, inspecting, launching and uninstalling of Java applications is referred to generally as *application management*. In typical desktop computing environments, these tasks can be performed by utilizing the facilities of the host operating system. However, the situation is very different in many small, resource-constrained devices which often lack even basic facilities such as a built-in file system.

To facilitate the porting of KVM to small, resource-constrained platforms, KVM implementation contains an optional component called *Java Application Manager* (JAM) that can be used as a starting point for machine-specific implementations.

Note – The JAM that is provided as part of the CLDC Reference Implementation is used primarily for compatibility testing purposes. This JAM implementation is not compatible with the requirements of J2ME profiles such as MIDP. To implement a MIDP-compliant Java Application Manager, refer to the MIDP Reference Implementation.

At the compilation level, JAM can be turned on or off by using the flag

```
#define USE_JAM 1
```

When building the KVM using `gnumake`, the following command automatically builds the system with the JAM enabled:

```
gnumake USE_JAM=true
```

This section provides a brief overview of the JAM reference implementation provided with KVM. The description below assumes that the target device has some kind of a “microbrowser” that can be used for initiating the downloading of applications. This microbrowser is commonly provided as part of the native computing environment, but it can also be part of the JAM in some implementations.

14.1 Using the JAM to install applications

Java Application Manager is a native C application that is responsible for downloading, installing, inspecting, launching, and uninstalling Java applications.

From the user's viewpoint, the JAM is typically used as follows:

1. The user sees an application advertised on a content provider's web page.
2. The user selects the tag to install it.
3. The Java application is downloaded and installed.
4. The user runs it.

Here's a more detailed description:

1. While browsing a content provider web page using a native microbrowser, the user sees a description of the Java application in the text of the page, and a highlighted tag (or button) that asks them if they want to install the application. The tag contains a reference to an application *Descriptor File*. The Descriptor File, typically with a .jam file extension, is a text file consisting of name/value pairs. The purpose of this file is to allow the JAM to decide, before it tries to download it, whether the Java application the user selected can be installed successfully on the device. This saves the user the cost of moving the Java application to the device if it cannot be installed. The Descriptor File is small (several hundred bytes), while a typical Java application is from 10 to 20 kilobytes, so it is much cheaper to download the Descriptor File rather than the entire Java application.
2. The user selects the tag to start the installation process. The browser retrieves the Descriptor File from the web site.
3. The browser transfers program control to the JAM, passing it the content of the Descriptor File and the URL for the page it was browsing.
4. The JAM checks to see if the application is already installed on the device, and checks its version number (see later discussion on the details of application updating.) It then reads the JAR-File-Size tag of the Java application to ensure that there is sufficient space on the device to save it.
5. If there is sufficient space to install the application, the JAM uses the JAR-File-URL tag in the descriptor file to get the URL of the JAR file (it may use the base URL to the Descriptor File, if the JAR-File-URL tag is a relative URL) and start the download process using HTTP. The JAM then stores the JAR file on the device.

If the download process is interrupted, the JAM discards the partially downloaded application, as if the application was never downloaded before.

6. The JAM adds the application to the list of installed Java applications, and registers it with any other native tools as required. The JAM saves the following information along with the JAR file:
 - name of JAR file,
 - absolute URL from where the JAR file was downloaded from,
 - main class of the java application,
 - name of the application,
 - version number of the application.

The absolute URL and the version number are used to uniquely identify an application during application update (see next subsection.)

In the reference implementation of the JAM, the user is shown the list of installed Java applications on the device, with the recently installed application selected for execution.

However, if the Use-Once tag is set to yes, JAM does not add the application to the list, and it launches the application immediately.

7. Any errors encountered during the process must be handled by the JAM. A help page URL for the content provider is included in the Descriptor File. The JAM can then direct the user to this URL using the native browser.

14.1.1 Application launching

Here's a typical use case for launching a Java application:

1. The user is shown a list of Java applications (the user interface design is left up to the manufacturer.)
2. The user selects the Java application that is to be launched (the user interface design and selection mechanism is left up to the manufacturer).
3. The JAM launches the KVM with a parameter containing the main class of the application. The KVM initializes the main class and starts executing it. As additional classes are required for the execution of the application, the KVM uses a manufacturer-defined API to unpack and load the class files from the stored JAR file.
4. The Java application is displayed on the screen to the user.
5. When the application exits, and if the Use-Once tag in the Descriptor File is set to YES, the JAM removes the JAR file.

14.1.2 Application updating

When the content provider updates an application (for example, to fix bugs or add new features), the content provider should do the following:

1. Assign a new version number to the application.
2. Change the Descriptor File of the application to use the new version number.
3. Post the updated JAR file on the content provider's web site, using the same JAR-File-URL tag as the previous version of the application.

When the user requests the installation of an application, the JAM checks if the application's JAR-File-URL is the same as one of the installed applications. If so, and the Application-Version of the requested version is newer than the installed version, the JAM prompts for user approval before downloading and installing the newer version of the application.

The reference implementation uses a string to specify the version number in the following format:

`Major.Minor[.Micro] (X.X[.X])`, where the `.Micro` portion is optional (it defaults to "0"). In addition, each portion of the version number is allowed to a maximum of 2 decimal digits (that is, the range is from 0 to 99.)

For example, "1.0.0" can be used to specify the first version of an application. For each portion of the version number, leading zeros are not significant. For example, "08" is equivalent to "8". Also, "1.0" is equivalent to "1.0.0". However, "1.1" is equivalent to "1.1.0", and not "1.0.1".

In the reference implementation, missing Application-Version tag is assumed to be "0.0.0", which means that any non-zero version number is considered as a newer version of the application.

The JAM must ensure that if the application update fails for any reason, the older version is left intact on the device. When the update is successful, the older version of the application is removed.

14.2 JAM components

14.2.1 Security requirements

The JAM, its data, and associated libraries, should be stored securely on the device. The device manufacturer must ensure that these components cannot be modified by Java applications or other downloadable content.

14.2.2 JAR file

JAR files are a standard feature of Java designed to hold class files and application resource data in a compressed format. JAM-compliant JAR files hold exactly one Java application and its associated resources. Compressed JAR files reduce the size of the application by approximately 40% to 50%. This both reduces the storage requirements on the device and reduces the download time for the application. Items in the JAR file are unpacked as required by the JAM.

14.2.3 Application Descriptor File

The Application Descriptor File is a readable text file. It consists of name-value pairs that describe the important aspects of its associated Java application. It is referenced from a tag on a content provider's web page. It is created and maintained by the Java application developer and stored along with its application JAR file on the same web site. Developers may create this file with any text editor.

The Descriptor File has the following entries (tag names are case sensitive):

`Application-Name`

Displayable text, limited to width of screen on the device

`Application-Version`

Major.Minor[.Micro] (X.X[.X], where X is a 1 or 2 digit decimal number, and the .Micro part is optional)

`KVM-Version`

Comma separated list of KVM version strings as defined in the CLDC microedition.configuration system property (see CLDC Specification). “CLDC-1.0” is an example of the KVM version string. The items in the list are matched against the KVM version string on the device, and an exact match is required to execute this application. Any item matching the KVM version string on the device satisfies this condition. For example, “CLDC-1.0, CLDC-1.0.3” runs on either version of KVM on the device.

Main-Class

Text name of the application's Main class in standard Java format.

JAR-File-Size

Integer in bytes

JAR-File-URL

Standard URL text format to specify the source URL. If this is a relative URL, then the URL to the Descriptor File is the base URL.

Use-Once

yes/no

Help-Page-URL

Standard URL text format, used by the browser to access help pages

Additional requirements and restrictions:

- The MIME type for the Descriptor File is `application/x-jam` and the extension is `.jam`.
- All URLs must point to the same server from which the web page was loaded.
- The JAM must store the Descriptor File contents, in a manufacturer-specific format for possible later use.

The application developer may add any application specific name-value pairs to the Descriptor File. This allows the application to be configured at deployment by changing the values in the Descriptor File. So, different Descriptor Files could use the same application JAR file, with different application parameters.

The format of the tag is a string, but it is recommended that it follow a similar style as the tags defined in the above table. The format of the value is an application specific string.

A simple proposed API to retrieve the value via the JAM could be:

```
public String GetApplicationParameter(String name)
```


14.2.4 Network communication

Whenever a Java application tries to make an HTTP connection, the networking implementation should check with the JAM to find the name of the server where the application was downloaded. This ensures that the connection is made to the same server the application came from. A string comparison is made between the host name in both the URLs.

14.3 Application lifecycle management

The lifecycle of a Java application is defined to be the following:

- The KVM task is launched and instructed to execute the main class of the Java application (as described by the Main-Class entry of the Descriptor File.)
- The Java application executes inside the context of the KVM task and responds to user events.
- The KVM task exits, either voluntarily, or involuntarily, and terminates the Java application.

The term task is used loosely to describe the KVM as a logically distinct execution unit. In actual devices, the KVM task can be implemented as a task, a process or a thread of the underlying operating system.

The API functions for controlling the lifecycle of the KVM are not specified, as the mechanism is vastly different from platform to platform. Instead, it is required that all JAM implementations support the following features:

- The JAM implementation must be able to launch the KVM task and start executing the `main` class of the Java application.
- The JAM implementation must be able to forcibly terminate the KVM task, and optionally be able to suspend and resume the KVM task.
- The suspension, resumption, and termination of the KVM must be performed by the procedures described below.

14.3.1 Termination of the KVM Task

The KVM task can be terminated in two ways: voluntarily or involuntarily.

The application can voluntarily terminate itself by calling the Java method `System.exit`. Under certain conditions, the JAM may decide to force the KVM to terminate. The exact method of triggering forced termination is platform dependent.

For example, the JAM may spawn a watchdog thread that wakes up after a certain period. If the watchdog thread detects that the KVM has not terminated voluntarily, it forces the KVM to terminate.

During forced termination, the JAM actively frees all resources allocated by the KVM and terminates the KVM task. The exact procedure is platform dependent. On some platforms, calling `exit` or `kill` may be enough. On other platforms, more elaborate clean-up may be required.

14.4 Error handling

The JAM is responsible for handling all errors encountered in installing and launching Java applications. The method of handling errors differs from implementation to implementation, but the JAM should be able to interact with the user to resolve the error if possible. To assist in this, the Descriptor File has a tag called `Help-Page-URL` that is set by the content provider. The JAM may decide that under certain conditions, the browser should be invoked and the user sent to the help page. The help page could have information that would allow the user to contact the content provider for additional assistance.

14.4.1 Error conditions

The following are a set of possible error conditions and sample messages that can be displayed to describe the error to the user. Manufacturers should design the messages so that they are appropriate to their device user interface.

- The user tries to install an application whose size is larger than the total storage space available on the device:

“NAMEOFAPP” is too large to run on this device and cannot be installed.

- The user tries to install an application, whose size is larger than the free storage space (but smaller than the total storage space) on the device:

There is not enough room to install. Try removing an application and trying again.

- The user tries to install an application that is already installed on the device.

“NAMEOFAPP” is already installed. (Soft buttons should be labeled OK and Launch. Launch would run the existing application on the device.)

- The user tries to install an application that is not designed for the particular device they own.

“NAMEOFAPP” won’t work on this device. Choose another application. (Soft button label = Back, Done.)

- The user tries to install an application and the tags describing the Java application have a syntax error or an invalid format that results in installation failure.

The installation failed. Contact your ISP for help.

- The user tries to install an application, the URL to the application is incorrect or inaccessible, and the application cannot be installed.

The URL for “NAMEOFAPP” is invalid. Contact your ISP for help.

- The user tries to install an application, the application is not the same size as described in the Descriptor File. The application should be discarded.

“NAMEOFAPP” does not match its description and may be invalid. Contact your ISP for help.

- The user is installing an application. During application download, the connection drops, and the application is not loaded onto the device successfully.

The connection dropped and the installation did not complete. Please try installing again. [Soft button label = Install, Back]

- The user is installing an application, and the URL specified matches exactly with the one located already on the device.

The JAM should check the version # of both versions and present a decision to the user.

- The user tries to run an application and for some reason the application cannot launch (for example, the JAM failed to create a new OS task to run the KVM).

Cannot launch “NAMEOFAPP”. Contact your ISP for help.

- The user has been running an application. The application tries to save to the scratchpad and fails.

Cannot save data. Contact your ISP for help.

- The user is running an application and it crashes or hangs during execution.
NOTE: This is a generic error.

“NAMEOFAPP” has unexpectedly quit.

Java-Level Debugging Support (KDWP)

Starting from release 1.0.2, KVM provides facilities for plugging the K Virtual Machine into a third-party Java development and debugging environment that is compliant with the *JPDA (Java Platform Debug Architecture)* specification supported by Java 2 Standard Edition. Further information on the JPDA architecture is available at <http://java.sun.com/products/jpda/>.

Due to strict memory constraints, KVM does not implement support for the JVMDI (Java Virtual Machine Debug Interface) and full JDWP (Java Debug Wire Protocol) specifications required by JPDA.

Instead, KVM implements a subset of the JDWP known as *KDWP (KVM Debug Wire Protocol)*. A specification of the KDWP protocol is available in a separate document listed in Section 1.2, “Related Documentation.”

Note – The Java-level debugging facilities provided in this 1.0.3 release have been tested under Sun’s *Forte* development environment (Community Edition 1.0 Update Release 2 and Forte for Java 2.0) and *JDB* (version 99/06/11). In addition, we have run tests with Metrowerks CodeWarrior 6.0 for Java and Borland JBuilder 4.0 Enterprise Edition. At this point, since the 3rd party development environments for J2ME are still under development, bugs may still be found. Any comments or feedback on the debugging interface would be highly appreciated. Also note that since KDWP implements a subset of the full Java Debug Wire Protocol, not all debugger commands can be supported. Refer to the *KDWP Specification* document for further details.

15.1 Overall architecture

The KDWP was designed to be a strict subset of the JDWP, primarily based on the resource constraints imposed on the KVM. In order to make KVM run with a JPDA-compatible debugger IDEs, a *debug agent* (debug proxy) program is interposed between the KVM and the JPDA-compatible debugger. The debug agent allows many of the memory-consuming components of a JPDA-compliant debugging environment to be located on the development workstation instead of the KVM, therefore reducing the memory overhead that the debugging interfaces have on the KVM and target devices. As obvious, the debugging interfaces can be turned off completely (at compile time) on those platforms/ports that do not need Java-level debugging support.

At the high level, the Java-level debugging support implementation consists of two parts:

- the actual code in the KVM to support a subset of the JDWP, and
- the debug agent that performs some of the debug commands on behalf of the KVM.

The overall architecture for the Java-level debugging interface is illustrated in Figure 15-1. In that figure, the topmost box represents the JPDA-compliant debugging environment (“JPDA Debugger”) running on a development workstation. The debugger is connected to the debug agent that talks to the KVM.




FIGURE 15-1 Java-level debugging interface architecture

The debug agent connects to the KVM via a socket connection. Similarly, the debugger connects to the debug agent over a socket. The debugger is unaware that it is connected to the debug agent. The debugger appears to be communicating directly

with a JDWP-compliant Java Virtual Machine. In fact, the debug agent can be configured in *pass through mode* so that all packets are passed from input to output using the debug agent with a standard Java VM. In normal KVM debug mode, the debug agent examines packets from the debugger and determines which packets are to be handled by the KVM and which are to be handled within the debug agent.

The main processing done in the debug agent is the parsing of class files to extract debugging information. This includes line number and code offset information and variable information. The KDWP implementation within the KVM includes some *vendor specific commands* that the debug agent uses to communicate with the KVM.

15.2 Debug Agent

The debug agent (also known as *debug proxy*) is written in the Java programming language and the code is in the KVM source tree under the directory `tools/kdp/src/kdp`. There are two main portions of the code: the portion that handles connections to the debugger and to KVM, and the portion that handles the parsing of the class files. The latter code is located in subdirectory `classparser`.

15.2.1 Connections between a debugger and the KVM

The portion of the code that handles connections to the debugger and to KVM resides in file `KVMDebugProxy.java`. This code creates two objects: *DebuggerListener* and *KVMListener*. The *DebuggerListener* class handles the retrieval of packets from the debugger, and the *KVMListener* class handles the retrieval of packets from the KVM. *DebuggerListener* and *KVMListener* are both subclasses of class *Thread*. Therefore, when they are invoked they start a new thread of execution (on the development workstation.) Each object also gets passed a handle to the other object (that is, the *KVMListener* object gets passed a handle to

the DebuggerListener object, and vice versa). This enables cross-communication of packets between the debugger and the KVM. The following diagram (Figure 15-2) may help to clarify this further:

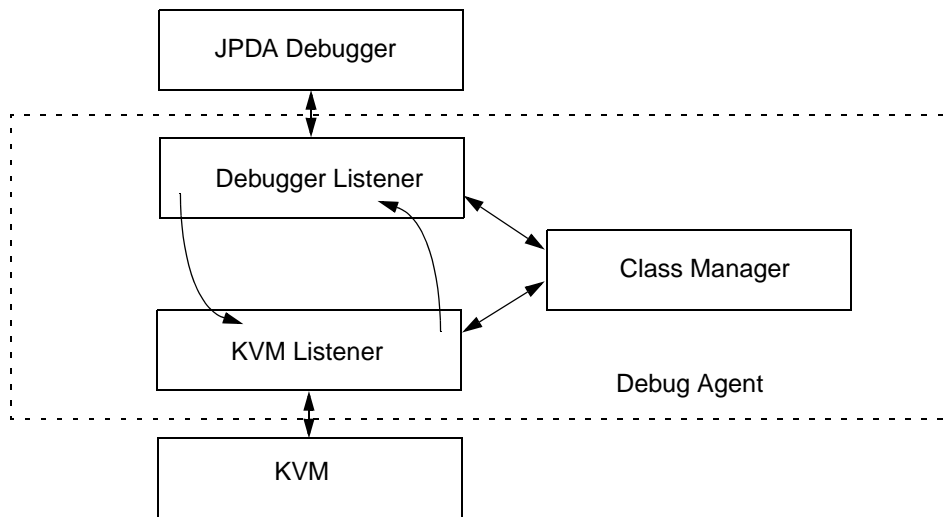


FIGURE 15-2 Debugger and KVM connections

In a typical scenario, the KVM is started with the `-debugger` flag, which puts it into a *debugger enabled* mode. In this mode the KVM listens on a socket for a connection from the debug agent. When the debug agent is started, it connects to this socket, and then listens on another socket for a connection from the debugger. When the debugger connects, it issues the *JDWP handshake* command, which consists of the string “JDWP-Handshake”. The debug agent acknowledges by reflecting this string back to the debugger. Meanwhile, the debug agent has sent the handshake command to the KVM and the KVM has responded back with information concerning which optional events it supports. The `KVMListener` then queries the KVM for a list of all the classes that are currently loaded into the VM. This information is used to build a hash table of `ClassFile` objects that is used later when the debugger requests information about a specific class (such as line number information, method information, etc.) At this point, each thread is listening for packets. The KVM sends a `VMInit` event to the debugger via the debug agent, which indicates to the debugger that the KVM is starting its execution of the Java application. The debugger might also send packets that indicate to the KVM to start up other events such as `ClassPrepare` or `ClassLoad`.

The communication code for the debug agent is in source file `SocketConnection.java`. In this file, each object (`KVMListener` and `DebuggerListener`) creates a thread of execution that waits for packets to arrive from its respective socket. If the packet is a command packet (the `Packet.Reply` bit is not set), then it puts that packet on a `packetQueue` list (see file

`ProxyListener.java`) and a notification is sent to any object waiting on that queue. The packet is then extracted from the queue by whatever listener is waiting for that packet on that queue. In the run method for the `KVMListener` and `DebuggerListener`, each packet is analyzed to determine if the debug agent needs to process the packet or whether it is to be transmitted to the other object for further processing.

15.2.2 Packet processing

The `DebuggerListener` object intercepts a number of packets as is evident by examining the code for the large switch statement located after the call to `waitForPacket`. When `waitForPacket` returns with a packet, the debug agent first creates a new `PacketStream` object, then checks to see if the debug agent needs to process that packet (For example, the `SENDVERSION_CMD` packet is processed by the debug agent directly, and a response is created and sent back to the debugger without any interaction with the KVM.) A more complex command would be the `FIELDS_CMD` of the `REFERENCE_TYPE_CMDSET`. For this command, the debugger has passed in a *class id*, which is used by the debug agent to find a `ClassFile` object via the `ClassManager.classMap` object. The `classMap` object is filled by the `KVMListener` object when it receives the `ClassPrepare` events from the KVM. Once the debug agent has obtained the `ClassFile` object, it uses the `getAllFieldInfo` method to obtain a *list* of fields, and iterates through this list passing the information back to the debugger. Once again, there is no interaction with the KVM.

Similarly, within the source file for the `KVMListener.java`, the `KVMListener` object intercepts the `CLASS_PREPARE` events that are passed up from the KVM. `KVMListener` creates a new `ClassFile` object via the call to `manager.findClass` and inserts it into the `ClassManager.classMap` hashtable. `KVMListener` then passes the event to the debugger so that it can process the event as well.

15.3 Debugger support within KVM

The debugger support within the KVM consists primarily of four source (`.c`) files under the `VmExtra/src` directory and three header (`.h`) files under `VmExtra/h` directory. All debugger code is included with the conditional compilation flag, `ENABLE_JAVA_DEBUGGER`. If this flag is enabled, and the KVM is rebuilt, then the Java debugger support is included within the KVM. If Java debugger support is not desired, set this define in `main.h` to 0.

Note – If your target platform or port does not require Java-level debugging support, we recommend turning the debugging code off at compile time (in file `main.h` or in your platform-specific `machine_md.h` file):

```
#define ENABLE_JAVA_DEBUGGER 0
```

This will make the KVM executable much smaller.

The primary file for the Java debugger support within the KVM is the source file `debugger.c`. This file contains all the support needed for the KDWP API. Socket communication is handled by the code in file `debuggerSocketIO.c`. The `debuggerInputStream.c` and `debuggerOutputStream.c` files contain the code for handling the transmission of data being sent to/from the debugger support functions in `debugger.c`. The code in `debugger.c` file services all the KDWP requests that are sent by or through the debug agent. The function `processDebugCmds` handles the parsing of input packets to determine which command set and what command within the command set the packet is referring to. This function then determines the appropriate function that is to be invoked for handling this command. The `inputStream` handle as well as the `outputStream` handles are passed as parameters, and used for handling the reply back to the debug agent. For performance reasons, most commands use a global `inputStream` and `outputStream`. If these are already in use, another one is allocated from the heap.

15.3.1 Events

Events are essentially commands generated by the KVM. Events are passed up to the debug agent, which may in turn pass them up to the debugger. The code for handling an event will appear as follows:

```
#if ENABLE_JAVA_DEBUGGER

{
    CEModPtr cep = GetCEModifier();

    cep->thread = thisThread;

    setEvent_ThreadStart(cep);

    FreeCEModifier(cep);
}

#endif /* ENABLE_JAVA_DEBUGGER */
```

This creates a new `CEModPtr` structure that contains state information for this particular event. It then invokes a routine in `debugger.c`, which attempts to send the event. A typical event routine in `debugger.c` first determines if the event attempting to be sent has been enabled by a previous *Set Event* command from the debugger (via the `checkNOTIFY_WANTED` macro.) Then, `findSatisfyingEvent` is invoked to determine if this particular event matches an event request sent down from the debugger. The `findSatisfyingEvent` function also checks the event counter as well as any modifiers that the debugger has applied to this event. If the event passes, then it is sent on the `outputStream`. After an event is sent, `handleSuspendPolicy` is invoked to process whatever suspend policy the debugger has attached to this event when the debugger had issued the *Set Event* command. Some events such as breakpoints or single stepping will generally have a suspend policy of `ALL`, which means that all threads are suspended and that the KVM will essentially spin through the reschedule loop at the top of the interpreter loop waiting for a thread to resume. The *Resume* command will eventually come from the debugger when the user issues a *Continue* command or when the user explicitly issues a *Resume Thread* command.

In certain situations, events need to be deferred. This is because it is not possible to send the event to the debug agent and subsequently suspend the KVM threads, since the interpreter might be in the midst of executing a byte code. Thus in such cases, `insertDebugEvent(cep)` is called instead of `setEvent_XXX(cep)`, as shown in the example above. At the top of the interpreter loop (see `VmCommon/src/execute.c`), the events are checked, and if there is a pending event, it is sent when it is safe to do so.

15.3.2 Breakpoints

When a *Set Event* command is received to add a breakpoint, the code for handling the breakpoint event determines if the opcode at that particular location is a *Fast opcode*. If so, then the original opcode must be retrieved from the *inline cache* before the breakpoint is added. The original opcode is stored in an `EVENTMODIFIER` structure that is pointed to by the `VMEvent` structure for this particular event. When the Java bytecode interpreter hits the Breakpoint opcode (see `bytecodes.c`), and if not *single stepping*, then the `handleBreakpoint` function is invoked. This function restores the original opcode into the `thread` structure for the `CurrentThread`, at the point where the breakpoint had been entered. It then also sends an event to the debugger via the debug agent. Eventually, the user will press the *Continue* button on the debugger, which results in all threads to resume execution. The `RESCHEDULE` macro (see `execute.h`) includes some code in it for determining if this thread was just at a breakpoint, and if so, it will retrieve the next bytecode from a known location within the `thread` structure. The code within the interpreter loop will then execute this instruction.

15.3.3 Single stepping

When the debugger issues a `SingleStep` event request, the code in `debugger.c` must determine which type of step function it is (that is, *step by bytecode* or *step by line*), whether the step is a *Step Into* (step into a function), *Step Over* (step over calls to functions; that is, do not single step into another function), or *Step Out* (go back to the function that called this function). Additionally, if it is a *step by line*, then KVM needs to know what the code offset is for the next line number. To obtain this information, KVM calls a private API within the debug agent to return the target offset and the next line offset. The debug agent returns this information back to the KVM, which stores it into a `stepInfo` structure, which is part of the `threadQueue` structure (see `thread.h`.) Within the interpreter loop, a flag is checked to determine if this particular thread is in single step mode. If so, then the `handleSingleStep` function in `debugger.c` is invoked to process this *single step*. The `handleSingleStep` function determines if the instruction pointer has reached the target offset or if it has popped up a frame or if it has gone beyond the target offset. Depending on the type of stepping being performed, this function will determine when to send a `SingleStep` event to the debugger. In most cases, if the user is single stepping line by line, and when the code offset is equal to the target offset, it results in a `SingleStep` event to be sent to the debugger. All threads are typically suspended at this point, and as was the case for the breakpoint scenario above, the KVM will wait until the debugger resumes the threads via a `Continue` command or a subsequent `SingleStep` event.

15.3.4 Suspend and nosuspend options

It is desirable in certain IDE environments such as Borland's JBuilder to provide an option similar to that available in J2SE for starting up the KVM in two different debugging modes. Thus, as of KVM 1.0.3, the KVM debugger can be started in the following two modes:

```
kvm -debugger -suspend ...
```

```
kvm -debugger -nosuspend ...
```

In the `"-suspend"` mode (this is the default), the KVM will stop all the Java threads upon VM startup and wait for further commands from the IDE (development and debugging environment) before the debugging session proceeds any further. In the `nosuspend` mode, the Java threads start running immediately when the KVM is started.

In the most common cases, the KVM debugger is usually invoked in the `suspend` mode. Unless the application program being debugged requires substantial processing or is recursive in nature, it may not make much sense to invoke the KVM

in the *nosuspend* mode. This is because it is very likely that a simple application program may complete execution long before the debugger IDE is able to issue any commands to the KVM.

15.4 Using the Debug Agent and the JPDA Debugger

In order to run the debug agent, it is necessary to build the application class or classes being debugged to include debug information. It is also necessary to transform the application class file(s) using the preverifier. Then, after the KVM is invoked on a specified host and port, the debug agent can be started such that it listens to KVM requests on the KVM port, and a local port is specified for connecting with a JPDA-compatible debugger.

The following section summarizes the steps necessary to start a debug session in much more detail.

Note – KVM debugger functionality is already integrated into the J2ME Wireless Toolkit (WTK) 1.0.3. Therefore, if you are running WTK 1.0.3, the detailed steps in the next section are not necessary.

15.4.1 Starting a debug session

To start a debug session, the following five steps are necessary:

1. **Build the application classes to be debugged with the `-g` option to include debug information. Then, place the output in a separate directory for transforming the resulting class file. See Chapter 12, “Class File Verification.”**

```
javac -g -classpath <path> -d <directory> <class>
```

- `-g` indicates to include debug information
- `-classpath <path>`, where `<path>` indicates the directory in which the CLDC/KVM Java library classes and the application classes for the application being debugged are located.
- `-d <directory>`, where `<directory>` indicates the directory in which output classes will be written. The default output directory is `./output`.
- `<class>` is the application class or classes being debugged.

2. Invoke the preverifier for transforming the class file.

```
preverify -classpath <path> -d . <directory>
```

This will transform all classes under <directory> and places the transformed class files in the current directory (as specified by the -d option).

3. Start the KVM process:

```
kvm -debugger -classpath <path> -port <KVM port> <class>
```

- -debugger indicates to put the KVM in debugger enabled mode
- -classpath <path>, where <path> specifies the directory in which the CLDC/KVM Java library classes as well as the application classes for the application being debugged are located.
- -port <KVM port> is the KVM port. The default KVM port is 2800. This must match the KVM port specified by the debug agent below.
- <class> is the application class being debugged.

4. Start the debug agent (debug proxy):

```
java -classpath <path> kdp.KVMDebugProxy -l <localport> -p -r  
<KVM host> <KVM port> -cp <KVM_path>
```

- -classpath <path>, where <path> specifies the directories in which the debug proxy classes are located.
- -l <localport>, where <localport> is the port that the debugger connects to.
- -p indicates to use the class parser.
- -r <KVM host>, where <KVM host> is the remote host name.
- <KVM port> is the KVM port. As stated earlier, this port must match the KVM port specified in step 3 above.
- -cp <path>, where <path> is the directory or directories where the CLDC/KVM Java library classes as well as the application classes for the application being debugged are located.

5. Connect to the debug agent with the debugger:

- For Forte debugger, go to the *Debug->Connect* dialog box and insert the host where the debug agent is running and the local port number that had been specified using the -l <localport> option.

Note – To download the Forte debugger or for further information on Forte, please refer to the Forte tools website at

<http://www.sun.com/forte/ffj/index.html>. When running the Forte debugger, JDK 1.3 must be previously installed and be on the classpath, since only this version (or later) of the JDK includes support for the JPDA. For further information on downloading the JDK 1.3, please refer to the website at <http://java.sun.com/j2se/1.3/jre/>.

- For jdb (Java debugger), the command will be as follows:

```
jdb -attach <agent hostname>:<localport>
```

15.4.2 Debugging example

If the KVM is running on a system called *sicily*, and the debug agent and debugger are running on *debughost*, then the commands for starting the debug session would appear as follows:

- On the *sicily* system, build the application test as follows:

```
javac -g -classpath ../api/classes:../samples/classes  
-d output test.java
```

- Invoke the preverifier for building a preverified class file.

```
preverify -classpath ../api/classes:../samples/classes  
-d . output
```

- On the *sicily* system, type the following command to invoke the KVM:

```
kvm -debugger -classpath ../api/classes:../samples/classes  
-port 2800 test
```

- On the *debughost* system, assuming the current directory is *tools/kdp/classes*, then the following command would invoke the debug agent (debug proxy).

```
java -classpath . kdp.KVMDebugProxy -l 1234 -p -r sicily 2800  
-cp ../../../../api/classes:../../../../samples/classes
```

- Invoke the Forte debugger.

Select the *Debug->Connect* dialog and select the *socket transport*. Then, enter *debughost* in the hostname box, and enter 1234 in the port number box. Then, press OK.

