



CDC Runtime Guide

*for the Sun Java Connected Device Configuration
Application Management System*

Version 1.0

Sun Microsystems, Inc.
www.sun.com

November 2005

Copyright © 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.sun.com/patents> and one or more additional patents or pending patent applications in the U.S. and in other countries.

THIS PRODUCT CONTAINS CONFIDENTIAL INFORMATION AND TRADE SECRETS OF SUN MICROSYSTEMS, INC. USE, DISCLOSURE OR REPRODUCTION IS PROHIBITED WITHOUT THE PRIOR EXPRESS WRITTEN PERMISSION OF SUN MICROSYSTEMS, INC.

U.S. Government Rights - Commercial software. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements. This distribution may include materials developed by third parties.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and in other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, J2ME, Java ME, Sun Corporate Logo and Java Logo are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Products covered by and information contained in this service manual are controlled by U.S. Export Control laws and may be subject to the export or import laws in other countries. Nuclear, missile, chemical biological weapons or nuclear maritime end uses or end users, whether direct or indirect, are strictly prohibited. Export or reexport to countries subject to U.S. embargo or to entities identified on U.S. export exclusion lists, including, but not limited to, the denied persons and specially designated nationals lists is strictly prohibited.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright © 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. détient les droits de propriété intellectuels relatifs à la technologie incorporée dans le produit qui est décrit dans ce document. En particulier, et ce sans limitation, ces droits de propriété intellectuelle peuvent inclure un ou plus des brevets américains listés à l'adresse <http://www.sun.com/patents> et un ou les brevets supplémentaires ou les applications de brevet en attente aux Etats - Unis et dans les autres pays.

CE PRODUIT CONTIENT DES INFORMATIONS CONFIDENTIELLES ET DES SECRETS COMMERCIAUX DE SUN MICROSYSTEMS, INC. SON UTILISATION, SA DIVULGATION ET SA REPRODUCTION SONT INTERDITES SANS L'AUTORISATION EXPRESSE, ECRITE ET PREALABLE DE SUN MICROSYSTEMS, INC.

Cette distribution peut comprendre des composants développés par des tierces parties.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, J2ME, Java ME, Sun Corporate Logo et Java Logo sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Les produits qui font l'objet de ce manuel d'entretien et les informations qu'il contient sont régis par la législation américaine en matière de contrôle des exportations et peuvent être soumis au droit d'autres pays dans le domaine des exportations et importations. Les utilisations finales, ou utilisateurs finaux, pour des armes nucléaires, des missiles, des armes biologiques et chimiques ou du nucléaire maritime, directement ou indirectement, sont strictement interdites. Les exportations ou reexportations vers des pays sous embargo des Etats-Unis, ou vers des entités figurant sur les listes d'exclusion d'exportation américaines, y compris, mais de manière non exclusive, la liste de personnes qui font l'objet d'un ordre de ne pas participer, d'une façon directe ou indirecte, aux exportations des produits ou des services qui sont régis par la législation américaine en matière de contrôle des exportations et la liste de ressortissants spécifiquement désignés, sont rigoureusement interdites.

LA DOCUMENTATION EST FOURNIE "EN L'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFACON.

Contents

Preface xiii

1. Introduction 1–1

- 1.1 Goals 1–2
- 1.2 Usage Contexts 1–2
- 1.3 CDC Technology Implementations 1–3
- 1.4 CDC Target Device Requirements 1–4
- 1.5 Java ME Technology Standards 1–4
- 1.6 Java ME API Choices 1–6
- 1.7 CDC Application Features 1–7
- 1.8 Developer Tools 1–7
- 1.9 Application Management 1–8

2. Installation 2–1

- 2.1 Installing the CDC Java Runtime Environment 2–2
 - 2.1.1 Local Installation 2–2
 - 2.1.2 Remote Installation 2–2
 - 2.1.3 Installing CDC AMS on a Zaurus personal mobile tool 2–3
 - 2.1.4 Installing CDC AMS on a Cobalt Qube 2–4
- 2.2 Installing an Optional Package 2–4

2.3	Testing the CDC Java Runtime Environment	2-4
2.4	Removing the CDC Java Runtime Environment	2-5
3.	Software Contents	3-1
3.1	Standard Files	3-1
3.2	CDC AMS Files	3-3
3.3	Security Files	3-5
3.4	Development Files	3-6
3.5	Test and Demonstration Files	3-7
4.	Running Applications	4-1
4.1	Launching a Java Application	4-1
4.2	Running Managed Applications (<i>Personal Basis Profile and Personal Profile only</i>)	4-2
4.2.1	Running an Applet (<i>Personal Profile only</i>)	4-2
4.2.2	Running an Xlet (<i>Personal Basis Profile and Personal Profile only</i>)	4-3
4.3	Class Search Path Basics	4-4
4.3.1	Java Class Search Path	4-5
4.3.2	Native Method Search Path	4-6
4.4	Memory Management	4-7
4.4.1	The Java Heap	4-7
4.4.2	Garbage Collection	4-8
4.4.2.1	Garbage Collection in the CDC HotSpot Implementation	4-8
4.4.2.2	Default Generational Collector	4-9
4.4.2.3	Tuning Options	4-11
4.4.3	Class Preloading	4-11
	Class Preloading and Verification	4-12
4.4.4	Setting the Maximum Working Memory for the Dynamic Compiler	4-13

4.5	Tuning Dynamic Compiler Performance	4-13
4.5.1	Dynamic Compiler Overview	4-14
4.5.2	Dynamic Compiler Policies	4-15
4.5.2.1	Managing the Popularity Threshold	4-16
4.5.2.2	Managing Compiled Code Quality	4-17
4.5.2.3	Managing the Code Cache	4-17
5.	Application Management	5-1
5.1	Launching CDC AMS	5-1
5.2	Exploring CDC AMS	5-3
5.2.1	AwtPDA Presentation Mode (<i>Personal Profile only</i>)	5-4
5.2.2	PBP Presentation Mode (<i>Personal Basis Profile only</i>)	5-6
5.3	Launching Applications	5-6
5.3.1	AwtPDA Presentation Mode (<i>Personal Profile only</i>)	5-7
5.3.2	PBP Presentation Mode (<i>Personal Basis Profile only</i>)	5-8
5.4	Switching Applications	5-8
5.5	Terminating Applications	5-10
5.6	Installing Applications	5-10
5.6.1	Manual Installation	5-11
5.6.2	OTA Provisioning	5-11
6.	Security	6-1
6.1	Overview	6-2
6.1.1	Built-in Security Features	6-2
6.1.2	Security Policy Framework	6-2
6.1.3	Security Provider Architecture	6-3
6.2	Security Procedures	6-4
6.2.1	Using Alternate Security Providers	6-4
6.2.2	Public Key Management	6-4

6.2.3	Security Policy Management	6-5
6.2.4	Seed Generation for Random Number Generation	6-5
7.	Localization	7-1
7.1	Setting Locale System Properties	7-1
7.2	Timezone Information Files	7-2
7.3	Font Management (<i>Personal Basis Profile and Personal Profile only</i>)	7-2
8.	Developer Tools	8-1
8.1	Compiling With <code>javac</code>	8-1
8.1.1	CDC and Java SE	8-1
8.1.2	Compiling Java Source Code for the Java SE Platform	8-3
8.1.3	Compiling Java Source Code for CDC	8-3
8.1.4	Determining the Target Class Library	8-4
8.1.5	Useful <code>javac</code> Command-Line Options	8-4
8.1.5.1	<code>-classpath</code> <i>classpath</i>	8-5
8.1.5.2	<code>-bootclasspath</code> <i>classpath</i>	8-5
8.1.5.3	<code>-extdirs</code> <i>classpath</i>	8-5
8.1.5.4	<code>-source</code> <i>release</i>	8-5
8.1.5.5	<code>-target</code> <i>version</i>	8-5
8.1.5.6	<code>-deprecation</code>	8-5
8.1.6	Compiling an Example CDC Program	8-6
8.2	Debugging With <code>jdb</code>	8-6
8.2.1	Debug Command-Line Options	8-7
8.2.2	Running the Debug Version of <code>cvm</code>	8-7
8.2.3	Running <code>jdb</code> on the Host Development System	8-8
8.3	Profiling with <code>hprof</code>	8-9
8.3.1	Profiling Command-Line Options	8-9
8.3.2	Running <code>cvm</code> With <code>hprof</code>	8-10

A. cvm Reference A-1

- A.1 Synopsis A-1
- A.2 Description A-1
- A.3 Options A-2

B. mtask Command Language Reference B-1

- B.1 Launch Command Language B-2
- B.2 Warming Up the mtask Server B-3

C. cvmc Reference C-1

- C.1 Synopsis C-1
- C.2 Description C-1
- C.3 Options C-1

D. Java ME System Properties D-1**E. Serial Port Configuration Notes E-1**

- E.1 Serial Port Setup E-2
- E.2 OS-Level Testing E-3

F. Platform Font Administration Notes F-1**G. Zaurus Installation Notes G-1**

- G.1 Zaurus System Requirements G-1
- G.2 Installation Procedure G-2

H. Cobalt Installation Notes H-1

- H.1 Installation Procedure H-1

I. Provisioning Server Notes I-1

- I.1 Download the J2EE 1.3.1 SDK I-1
- I.2 J2EE Server Setup I-2

- I.3 Download the CPRI Server I-3
- I.4 CPRI Server Setup I-3
- I.5 Deploy a PAR File I-4
- I.6 Testing I-6

Figures

FIGURE 1-1	CDC Target Device Categories	1-1
FIGURE 1-2	An Example CDC Java Runtime Environment	1-6
FIGURE 4-1	GC Generations	4-10
FIGURE 4-2	Interpreter-Based Method Execution	4-14
FIGURE 4-3	Compiling a Method	4-14
FIGURE 4-4	Executing a Compiled Method	4-15
FIGURE 5-1	Launch an <code>mtask</code> Server JVM Instance	5-2
FIGURE 5-2	<code>cvmc</code> Sends an <code>mtask</code> Command to Clone the Server JVM Instance	5-2
FIGURE 5-3	AppManager Sends an <code>mtask</code> Command to the Server JVM Instance	5-2
FIGURE 5-4	Zaurus Qtopia Application Environment	5-3
FIGURE 5-5	AwtPDA Presentation Mode Login Screen	5-4
FIGURE 5-6	AwtPDA Presentation Mode Application Management Screen	5-5
FIGURE 5-7	PBP Presentation Mode Application Management Screen	5-6
FIGURE 5-8	Phone Xlet	5-7
FIGURE 5-9	CD Player Xlet	5-8
FIGURE 5-10	Taskbar Utility	5-9
FIGURE 6-1	Java Security Policy Model	6-3
FIGURE 8-1	CDC and Java SE API Compatibility	8-2
FIGURE 8-2	Compiling Java Source Code for the Java SE Platform	8-3
FIGURE 8-3	Compiling Java Source Code for CDC	8-4

[**FIGURE I-1**](#) ri-test Page I-4

[**FIGURE I-2**](#) ri-test Upload Page I-5

Tables

TABLE 1-1	Java ME API Standards 1–5
TABLE 2-1	Target Platforms 2–1
TABLE 2-2	Binary Distribution Archives 2–2
TABLE 3-1	Standard Files 3–1
TABLE 3-2	CDC AMS Files 3–3
TABLE 3-3	Security Files 3–5
TABLE 3-4	Development Files 3–6
TABLE 3-5	Test and Demonstration Files 3–7
TABLE 4-1	XletRunner Command-Line Options 4–3
TABLE 4-2	GC Runtime Options 4–11
TABLE 5-1	Application Repository Sub-Directories 5–11
TABLE 6-1	Security Documentation for the Java SE Platform 6–1
TABLE 7-1	Locale System Properties 7–1
TABLE 7-2	Font Management Comparison 7–2
TABLE 7-3	Logical Font Names 7–3
TABLE 8-1	cvm Debugging Options 8–7
TABLE 8-2	hprof Command-Line Options 8–9
TABLE A-1	Java SE Command-Line Options A–2
TABLE A-2	CDC-Specific Command-Line Options A–5
TABLE A-3	<i>-Xgc:suboption</i> A–7

TABLE A-4	-Xopt : <i>suboption</i> A-7
TABLE A-5	-Xtrace : <i>flags (Ol only, unsupported)</i> A-7
TABLE A-6	JVMDI Options A-8
TABLE A-7	JVMPI Options A-9
TABLE A-8	-Xjit : <i>options (Ol only)</i> A-9
TABLE A-9	-Xjit : <i>inline=suboption (Ol only)</i> A-11
TABLE A-10	-Xjit : <i>compile=suboption (Ol only)</i> A-11
TABLE A-11	-Xserver : <i>suboption (CDC AMS only)</i> A-11
TABLE A-12	-Xjit : <i>trace=option (Ol only)</i> A-12
TABLE B-1	Launch Commands B-2
TABLE C-1	cvmc Command-Line Options C-2
TABLE D-1	CDC System Properties D-1
TABLE E-1	Serial Communications References E-1
TABLE F-1	Font Management References F-1
TABLE G-1	Zaurus System Requirements G-1
TABLE G-2	Useful Zaurus Resources G-2

Preface

This runtime guide describes how to use a Java runtime environment based on the Connected Device Configuration (CDC) with its related profiles and optional packages. It focuses on runtime issues like deployment, configuration and running application software based on Java technology, as well as developer issues like compiling, debugging and profiling.

This runtime guide is based on a version of the CDC Java runtime environment that has been tested and used in a development environment on test devices. So the information in this runtime guide is not what a typical end-user would generally need for two reasons:

- This runtime guide doesn't describe a specific end-user product implementation of CDC technology. [Chapter 1](#) shows how these product implementations can vary depending on the target device and the included optional APIs. So the user experience for product devices is based on decisions made by product designers who adapt CDC technology to their products specific needs.
- This runtime guide is intended for use within a product development context, including both runtime and application development. From a developer's perspective, runtime issues generally exercise configuration, testing or debugging features of the CDC Java runtime environment.

The companion document *CDC Build System Guide* describes how to build a CDC Java runtime environment for a specific target device, including the build-time options that control functionality, testing and performance features. This runtime guide focuses on how to use those features at runtime.

Who Should Read This Runtime Guide

This runtime guide is intended for software engineers who need to work with a CDC Java runtime environment for one of the following purposes:

- Testing a CDC Java runtime environment
- Developing applications
- System integration
- Porting the CDC Java runtime environment
- Porting one of the CDC profiles or optional packages

How This Book Is Organized

- [Chapter 1](#) introduces the CDC platform, including its standards, target devices, application characteristics and developer tools.
- [Chapter 2](#) shows how to install the CDC Java runtime environment on a sample target device.
- [Chapter 3](#) describes the contents of a CDC Java runtime environment.
- [Chapter 4](#) shows how to launch and use application software based on Java technology with a CDC Java runtime environment.
- [Chapter 5](#) describes the application management system (AMS) for launching, controlling applications.
- [Chapter 6](#) describes security features and how they are related to the security framework provided by the Java Platform, Standard Edition (Java SE).
- [Chapter 7](#) describes localization procedures including font management, locale-specific system properties and timezone information files.
- [Chapter 8](#) shows how to integrate the CDC Java runtime environment with Java SE developer tools like `javac`, `jdb` and `hprof`.
- [Appendix A](#) describes the command-line options for the `cvm` application launch tool.
- [Appendix B](#) describes the `mtask` command language.
- [Appendix C](#) describes the `cvmc` driver utility.
- [Appendix D](#) describes system properties for the Java Platform, Micro Edition (Java ME). These include CDC-specific system properties.
- [Appendix F](#) describes platform-specific font administration for Linux, Qt and X11.
- [Appendix G](#) describes platform-specific installation procedures for the Zaurus personal mobile tool.
- [Appendix I](#) describes how to set up a provisioning server using the J2EE Client Provisioning RI server.

Typographic Conventions

TABLE P-1 Typographic Conventions

Typeface	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Command-line variable; replace with a real name or value	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be superuser to do this. To delete a file, type <code>rm filename</code> .

Runtime Documentation for the Java Platform Standard Edition

Because CDC is heavily based on Java Platform Standard Edition, it's important to be familiar with the documentation for Java Platform Standard Edition. TABLE P-2 describes the main web pages for the runtime documentation for Java Platform Standard Edition.

TABLE P-2 Java Standard Edition Runtime Documentation

URL	Description
http://java.sun.com/docs/index.html	Main documentation web page for the Java SE platform.
http://java.sun.com/j2se/1.4.2/relnotes.html	Release notes for the Java SE platform, version 1.4.2.
http://java.sun.com/j2se/1.4.2/docs/tooldocs/tools.html	Tool documentation for the Java SE platform, version 1.4.2.

Related Documentation

TABLE P-3 Related Documentation

Title	Description
<i>CDC: Java Platform Technology for Connected Devices</i>	This white paper introduces CDC and the devices and applications it supports.
<i>CDC Build System Guide</i>	CDC build system installation, configuration and testing.
<i>CDC Porting Guide</i>	Procedures and interface definitions for porting CDC, including its Java virtual machine and Java class library to an alternate target platform.
<i>CDC HotSpot Implementation Dynamic Compiler Architecture Guide</i>	Internals reference for the CDC HotSpot Implementation dynamic compiler.
<ul style="list-style-type: none">• <i>CDC Technology Compatibility Kit User's Guide</i>• <i>Foundation Profile Technology Compatibility Kit User's Guide</i>• <i>Personal Basis Profile Technology Compatibility Kit User's Guide</i>• <i>Personal Basis Profile Technology Compatibility Kit User's Guide</i>• <i>Security Optional Package Technology Compatibility Kit User's Guide</i>	User documentation for running the TCK validation suites.
<i>Java Language Specification</i>	<i>Java Language Specification</i> defines the Java programming language. See http://java.sun.com/docs/books/jls .
<i>Java Virtual Machine Specification</i>	Defines the Java class format and the virtual machine semantics for class loading, which are the basis for the operation of the Java runtime environment and its ability to execute Java application software on a variety of different target platforms. See http://java.sun.com/docs/books/vmspec .
<i>Java Native Interface (JNI)</i>	<i>The Java Native Interface: Programmer's Guide and Specification</i> (Addison-Wesley, 1999) by Sheng Liang describes the native method interface used by the CDC HotSpot Implementation Java virtual machine. See http://java.sun.com/docs/books/jni .
<i>Java Virtual Machine Debugger Interface (JVMDI)</i>	Defines an interface that allows debugger tools like jdb and third-party debuggers to interact with a debugger-capable Java runtime environment. See http://java.sun.com/products/jpda/doc/jvmdi-spec.html .

TABLE P-3 Related Documentation (*Continued*)

Title	Description
<i>Java Virtual Machine Profiler Interface (JVMPi)</i>	Defines an interface that allows the hprof profiler to interact with a Java runtime environment to measure application behavior. See http://java.sun.com/j2se/1.4.2/docs/guide/jvmpijvmpih.html .
<i>J2ME Unified Emulator Interface Specification (UEI)</i>	Defines an interface that allows an external developer tool to control an emulator for the running Java ME applications.
<i>Inside Java 2 Platform Security</i>	Describes the Java security framework, including security architecture, deployment and customization. See http://java.sun.com/docs/books/security .

Sun Documentation Resources

Sun provides online documentation resources for developers and licensees.

TABLE P-4 Sun Documentation Resources

URL	Description
http://docs.sun.com	Sun product documentation
http://java.sun.com/j2me/docs	Java ME technical documentation
http://developer.java.sun.com	Java Developer Services
http://www.sun.com/software/jpe	Java Partner Engineering
http://java.net	An open community that facilitates Java technology collaboration.

Terminology

These terms related to the Java™ platform and Java™ technology are used throughout this manual.

Java technology level (Java level)

Java technology based (Java based)

class contained in a Java
class file (Java class)

Java programming
language profiler (Java profiler)

Java programming
language debugger (Java debugger)

thread in a Java virtual
machine representing a
Java programming
language thread (Java thread)

stack used by a Java
thread (Java thread stack)

application based on Java
technology (Java application)

source code written in the
Java programming
language (Java source code)

object based on Java
technology (Java object)

method in an object based
on Java technology (Java method)

field in an object based on
Java technology (Java field)

a named collection of
method definitions and
constant values based on
Java technology (Java interface)

a group of types based on
Java technology (Java package)

an organized collection of packages and types based on Java technology (Java namespace)

constructor method in an object based on Java technology (Java constructor)

exception based on Java technology (Java exception)

an application programming interface (API) based on Java technology (Java API)

a service providers interface (SPI) based on Java technology (Java API)

developer tool based on Java technology (Java developer tool)

system property in a Java runtime environment (Java system property)

security framework for the Java platform (Java security framework)

security architecture of the Java platform (Java security architecture)

Feedback

Sun welcomes your comments and suggestions on CDC technology. The best way to contact the development team is through the following e-mail alias:

cdc-comments@java.sun.com

You can send comments and suggestions regarding this runtime guide by sending email to:

docs@java.sun.com

Introduction

A Java runtime environment is an implementation of Java technology for a specific target platform. It performs a middleware function with features common to a native application: it is installed, launched and run like a native application. But its real purpose is to launch, run and manage Java application software on the target platform.

The Connected Device Configuration (CDC) Java runtime environment is an implementation of Java technology for connected devices. These include mobile devices like PDAs and smartphones as well as attached devices like set-top boxes, printers and kiosks.

CDC target devices can vary widely based on their features and purpose. **FIGURE 1-1** describes some CDC target device categories and organizes them by their two most important characteristics: purpose (fixed or general) and mobility (mobile or attached).

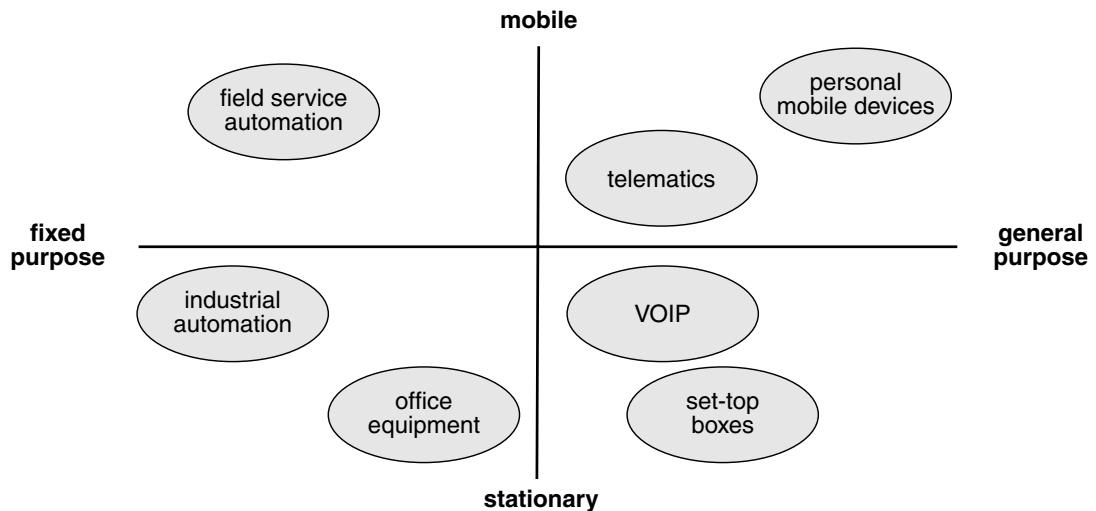


FIGURE 1-1 CDC Target Device Categories

This runtime guide describes how to use the CDC Java runtime environment for different purposes including application development, runtime development and solution deployment.

This chapter briefly introduces the CDC Java runtime environment through the following:

- [Goals](#)
- [Usage Contexts](#)
- [CDC Technology Implementations](#)
- [CDC Target Device Requirements](#)
- [Java ME Technology Standards](#)
- [Java ME API Choices](#)
- [CDC Application Features](#)
- Application Management
- [Developer Tools](#)

1.1 Goals

It is difficult to describe CDC technology without reference to the Java SE platform because Java SE represents the core of Java technology. In fact, the principal goal of CDC is to adapt Java SE technology from desktop systems to connected devices. Most of CDC's modifications to Java SE APIs are based on identifying features that are either too large or inappropriate for CDC target devices and then either removing or making them optional.

Other related goals of CDC include the following:

- Broaden the number of target devices for Java application software.
- Take advantage of target device features while fitting within their resource limitations.
- Provide a runtime implementation optimized for connected devices.
- Leverage Java SE developer tools, skills and technology.

1.2 Usage Contexts

The CDC Java runtime environment described in this runtime guide can operate in several different usage contexts:

- During *product development*, the CDC Java runtime environment has testing features that can help isolate problems while porting CDC technology to a new target platform. For example, the trace features provide details about opcode and method execution as well as garbage collection (GC) state.
- One of the final stages of product development is *TCK verification*. A TCK is a test suite that verifies the behavior of an implementation of Java technology. The TCK includes a test harness that runs a candidate Java runtime environment and launches a series of test Java applications. TCK verification is described in the TCK user guides listed in [“Related Documentation” on page xvi](#).
- *Application development* for the CDC platform requires a target Java class library for compiling Java source code and a CDC Java runtime environment for testing and debugging. [Chapter 8](#) provides more information about application development with the CDC Java runtime environment.
- When an application is complete and tested, it’s ready for *deployment*. CDC provides a number of deployment mechanisms including preloading with JavaCodeCompact, managed application models like applets and xlets and network-based provisioning systems.

1.3

CDC Technology Implementations

CDC technology is delivered by Sun through different kinds of software releases:

- A *Reference Implementation* (RI) demonstrates Java technology that is described in a *Java Specification Request* (JSR) and verified by a corresponding *Technology Compatibility Kit* (TCK). Because it serves a demonstration purpose, an RI does not provide the best available performance features.
- An *Optimized Implementation* (OI) is also a TCK-compliant implementation of Java technology. An OI provides the following benefits:
 - Undergoes more quality assurance (QA) testing
 - Provides superior performance
 - Supports a strategic platform or can be used as a starting point for porting Java technology to a different target platform

The sample OI described in this runtime guide is based on the Linux/ARM/Qt platform. It includes a dynamic compiler that provides the best available performance for Java applications. See [TABLE 2-1](#) for a description of the supported OI platforms.

1.4

CDC Target Device Requirements

CDC is an adaptable technology that can support a range of connected target devices that exist today and in the future. The baseline system requirements of these connected devices are the following:

- network connectivity
- 32-bit RISC-based microprocessor

The memory requirements for a CDC Java runtime environment vary based on the native platform, the profile and optional packages and the application. See [Section 4.4, “Memory Management” on page 4-7](#) for memory usage guidelines.

Other features of the CDC target device can include:

- a display for a graphical user interface (GUI)
- Unicode font support
- an open or proprietary native platform that provides operating system services

1.5

Java ME Technology Standards

CDC is part of the family of Java ME technology standards that support application software for connected devices. From an application developer’s perspective, CDC is a standards-based framework for creating and deploying application software on a broad range of consumer and embedded devices. The CDC APIs are largely based on well-known Java SE APIs, which makes the job of migrating skills, tools and source code easier. From a product designer’s perspective, CDC provides a standards-based Java runtime environment that supports a variety of target devices. This allows product designers to provide an application platform that fits within their device’s resource limitations while supporting a large number of applications and developers.

Java ME standards are developed in collaboration with industry leaders through the Java Community Process (www.jcp.org). JCP standards allow Java technology to adapt to the needs of evolving products in an open way by defining APIs that address common needs in application development. Furthermore, these standards allow product designers to choose which API features fit their product needs.

Java ME technology uses three kinds of API standards described in [TABLE 1-1](#) as building blocks that can be combined in a specific product solution.

TABLE 1-1 Java ME API Standards

Category	Description	Options
<i>Configuration</i>	Defines the most basic Java class library and Java virtual machine capabilities for a broad range of devices.	<ul style="list-style-type: none">• <i>Connected Device Configuration (CDC, JSR-218)</i> supports connected devices like smart phones, set-top boxes and office equipment.• <i>Connected Limited Device Configuration (CLDC, JSR-139)</i> supports small devices like cellphones.
<i>Profile</i>	Defines additional APIs that support a narrower range of devices. A profile is built on a specific configuration.	<ul style="list-style-type: none">• <i>Foundation Profile (JSR-219)</i> provides application-support classes like network and I/O support platforms without a standards-based GUI system.• <i>Personal Basis Profile (JSR-217)</i> provides a standards-based GUI framework for supporting lightweight components. In addition to the same application support classes provided by Foundation Profile, Personal Basis Profile includes support for the xlet application model.• <i>Personal Profile (JSR-216)</i> provides an AWT-based GUI toolkit. In addition to the same application support classes provided by both Foundation Profile and Personal Basis Profile, Personal Profile includes support for the applet application model.
<i>Optional Package</i>	Defines a set of technology-specific APIs.	<ul style="list-style-type: none">• The <i>Remote Method Invocation (RMI) Optional Package (JSR-66)</i> provides a subset of the Java SE RMI API for networked devices based on Java technology. It exposes distributed application protocols through Java interfaces, classes and method invocations and shields the developer from the details of network communications.• The <i>Java Database Connectivity (JDBC) Optional Package (JSR-169)</i> provides a subset of the JDBC 3.0 API that can be used by Java application software to access tabular data sources including spreadsheets, flat files and cross-DBMS connectivity to a wide range of SQL databases.• The <i>Java Secure Socket Extension (JSSE) Optional Package</i>, the <i>Java Cryptography Extension (JCE) Optional Package</i> and the <i>Java Authentication and Authorization Service (JAAS) Optional Package (JSR-219)</i> provide Java SE APIs for extending CDC's security architecture.• The <i>Web Services Optional Package (JSR-172)</i> provides standard access from Java ME clients to web services.

1.6

Java ME API Choices

Each Java ME licensee can create a Java runtime environment by choosing from a menu of standard APIs. The designer's choice must contain a configuration, a profile and any number of optional packages and these choices can vary from product to product. The critical point to understand is that the application developer must separately learn about which API combination are available for a specific CDC product implementation.

For example, [FIGURE 1-2](#) describes a Java runtime environment where a product designer selects CDC, Personal Profile, RMI Optional Package and JDBC Optional Package to represent a conforming CDC Java runtime environment.

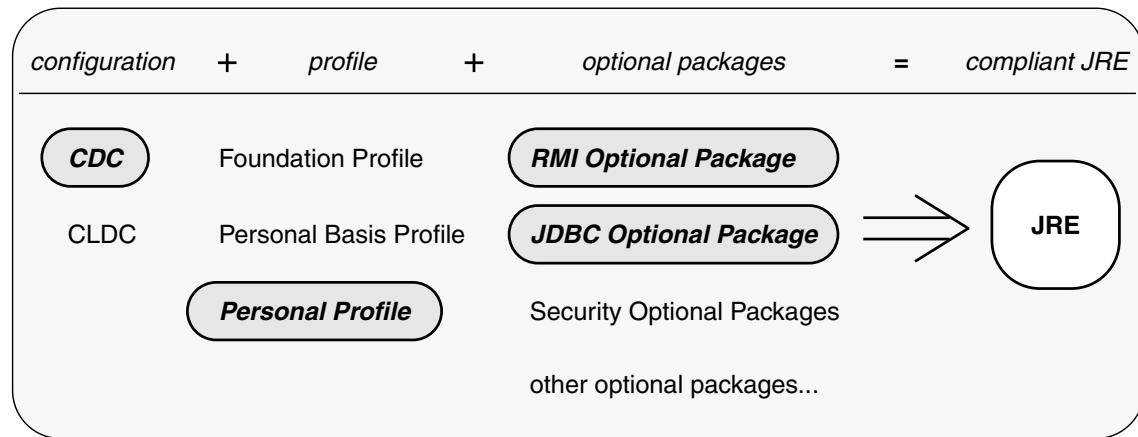


FIGURE 1-2 An Example CDC Java Runtime Environment

Note – See the companion document *CDC Build System Guide* for information on how to build a target development version of the CDC Java class library for application development that reflects the APIs chosen for a specific target product. [Chapter 8](#) describes how to compile Java application software with such a library.

1.7

CDC Application Features

The applications targeted by CDC technology have certain characteristics that distinguish them from the productivity tools and utilities common to desktop platforms.

- *Network connectivity.* The dominant trends in application development, like web browsers, XML-based web services and RSS, are based on network connectivity. Examples include the evolution of PDAs and cell phones into connected devices and the evolution of office printers into multi-function peripherals that can generate campus-specific reports.
- *Security.* Application developers and users are becoming increasingly aware of the need for security for their mobile and distributed applications. The Java SE security framework in CDC allows applications to use fine-grained security policies for application and enterprise security needs.
- *Application deployment.* Java technology has traditionally provided flexible application models. CDC profiles support managed application models like applets and xlets that allow developers to easily deploy applications over the network, either directly or through a provisioning server.
- *Standard data access.* Mobile clients need access to central databases to view and modify information. The JDBC and web Services optional packages provide standard data access for client-side applications.
- *Portable GUIs.* With the broad range of CDC target devices, applications need a GUI system that is flexible enough different user experiences and workflows while being portable enough to support different target devices. Personal Basis Profile and Personal Profile support conventional AWT-based GUIs as well as providing a hosting layer for building and supporting GUIs based on industry-standards and vendor-specific interfaces.

1.8

Developer Tools

Because CDC APIs are derived from Java SE APIs, application developers can migrate both their software and their skills to the CDC platform with little effort. Java SE developers can easily learn CDC APIs by focusing on their small differences with Java SE APIs. It is therefore easy to modify Java SE software for CDC devices. The ability to use Java SE developer tools like compilers, debuggers and profilers makes this transition easier.

The CDC Java runtime environment uses several developer tool-oriented specifications, including the following:

- Because CDC is based on the *Java Virtual Machine Specification* (see java.sun.com/docs/books/vmspec), application developers can use conventional Java SE compilers like javac.
- The *Java Virtual Machine Debugger Interface* (JVMDI, see java.sun.com/j2se/1.4.2/docs/guide/jvmdi) defines an interface that allows debugger tools like jdb and third-party debuggers to interact with a debugger-capable CDC Java runtime environment and explore the behavior of Java applications on a CDC target device.
- The *Java Virtual Machine Profiler Interface* (JVMPI, see java.sun.com/j2se/1.4.2/docs/guide/jvmdi) defines an interface that allows the hprof profiler to measure runtime data for a specific application or benchmark.
- The *J2ME Unified Emulator Interface Specification* (UEI) defines an interface that allows an external developer tool to control a Java ME emulator.
- cvm, the CDC application launcher, uses many command-line options that are available with java, the Java SE application launcher. Many of these options can be used for application testing and development.

Java SE tools like jar and keytool can also be used in CDC application development and deployment.

1.9 Application Management

The Sun Java Connected Device Configuration Application Management System (CDC AMS) is a process-based system for controlling multiple concurrent applications while sharing system resources in a robust and scalable manner. CDC AMS includes reusable components that can be adapted to build AMS implementations for specific products. To demonstrate these capabilities, CDC AMS includes two sample AMS implementations based on Personal Profile and Personal Basis Profile. This runtime guide describes how to use these sample AMS implementations to demonstrate different aspects of application management including OTA provisioning, application launching, application selection and application removal.

Installation

This chapter describes how to install and test a CDC Juntime environment on a sample target device. It shows how to perform the following procedures:

- Download and install a distribution bundle containing the CDC Java runtime environment.
- Perform a basic test procedure that demonstrates the operation of the CDC Java runtime environment.
- Install an optional package.
- Remove the CDC Java runtime environment.

TABLE 2-1 describes the target platforms for the CDC Java runtime environment.

TABLE 2-1 Target Platforms

Supported Platform	URL	Description
ARM-based Linux	http://www.myzaurus.com	The Sharp Zaurus SL-C860 is a Linux/ARM-based PDA.
MIPS-based Linux	http://www.linux-mips.org/wiki/index.php/Cobalt	The Cobalt Qube 2 is a Linux/MIPS-based network server.

2.1

Installing the CDC Java Runtime Environment

This section describes how to install a Java runtime environment contained in a Zip bundle described in [TABLE 2-2](#). The naming convention used is *api-1_1-release-bin-build-platform-CPU-date.zip*. The installation procedure described in [TABLE 2-2](#) uses the name *rt.zip* to represent one of these distribution archives.

TABLE 2-2 Binary Distribution Archives

Description	Archive File
Personal Basis Profile	appmgr_ppb-1_0-rr-bin-b21-linux-mips-15_nov_2005.zip
Personal Profile	appmgr_pp-1_0-rr-bin-b21-linux-arm-15_nov_2005.zip

The companion document *CDC Build System Guide* shows how to build a CDC Java runtime environment and prepare it for downloading onto a target device.

2.1.1

Local Installation

1. **Change the current directory to the location where you wish to install the CDC Java runtime environment.**

```
ftp> cd /home/test-cdc
```

2. **Download one of the binary distribution archives described in [TABLE 2-2](#).**

```
% get rt.zip
```

3. **Unzip the Zip bundle.**

```
% unzip rt.zip
```

2.1.2

Remote Installation

1. **Download one of the binary distribution archives described in [TABLE 2-2](#).**

2. Connect to the target system with an `ftp(1)` client.

This example is based on `ftp` and assumes that a directory named `/home/test-cdc` is available on the target device with enough space for installing and testing the CDC Java application environment.

```
% ftp cdc-dev  
Name (cdc-dev) :  
...  
Password:  
...
```

3. Change the current directory to the location where you wish to install the CDC Java runtime environment.

```
ftp> cd /home/test-cdc
```

4. Set the transfer type to binary.

```
ftp> binary
```

5. Transfer the Zip bundle to the target system.

```
ftp> put rt.zip
```

6. Quit the `ftp` client.

```
ftp> quit
```

7. Remotely login to the target system.

```
% ssh cdc-dev
```

8. Change the current directory to the directory containing the Zip bundle.

```
% cd /home/test-cdc
```

9. Unzip the Zip bundle.

```
% unzip rt.zip
```

After installation, the current directory should contain the contents of the CDC Java runtime environment. At a minimum, this includes two directories: `bin` and `lib` as well as two files: `democlasses.jar` or `testclasses.zip`. See [Chapter 3](#) for a complete description of the contents of a CDC Java runtime environment.

2.1.3 **Installing CDC AMS on a Zaurus personal mobile tool**

See [Appendix G](#) for platform-specific installation notes for the Zaurus personal mobile tool.

2.1.4 Installing CDC AMS on a Cobalt Qube

See [Appendix H](#) for platform-specific installation notes for the Cobalt Qube.

2.2 Installing an Optional Package

Optional packages can be installed with one of two methods:

- Integration into the CDC Java runtime environment at build-time. See the companion document *CDC Build System Guide* for more details.
- Installation into the optional package directory. See [Section 4.3, “Class Search Path Basics” on page 4-4](#) for a description of the extension class search path.

2.3 Testing the CDC Java Runtime Environment

You can quickly test the CDC Java runtime environment by running a sample application with `cvm`, the CDC Java application launcher:

```
% bin/cvm -cp testclasses.zip HelloWorld
Hello world.
```

If these test applications run successfully, the CDC Java runtime environment is installed correctly. See [Chapter 4](#) for more information about using `cvm` to run Java application software and [Appendix A](#) for a complete description of the command-line options.

Note – For the Personal Basis Profile or Personal Profile-based CDC Java runtime environment, there may be an extra step to specify the location of the Qt shared library on the Linux platform. See [Section 4.3.2, “Native Method Search Path” on page 4-6](#) for more information about the native library search path.

2.4

Removing the CDC Java Runtime Environment

The CDC Java runtime environment can be removed from the target system by removing the `bin` and `lib` directories.

```
% cd /home/test-cdc  
% rm -rf bin lib repository democlasses.jar testclasses.zip
```


Software Contents

A CDC Java runtime environment contains the software necessary to run Java applications on a target platform. The software contents of a CDC Java runtime environment can vary, especially during product development when different testing options may be selected at build-time. This chapter describes the organization of a CDC Java runtime environment, including standard files as well as optional security, developer and test files.

3.1 Standard Files

After installation, the CDC Java runtime environment is located in its installation directory. Because the location of this installation directory can be anywhere in the local file system, the CDC Java runtime environment specifies this location with the `java.home` system property. [TABLE 3-1](#) describes the standard files located in the installation directory based on the default build options.

TABLE 3-1 Standard Files

File	Description
<code>bin/cvm</code>	The CDC Java application launcher loads and executes Java applications.
<code>lib/class-lib.jar</code>	<i>Optional.</i> The CDC Java class library is used by the CDC Java runtime environment to locate and load core Java classes. The actual name of the archive file indicates the supported CDC specifications, e.g. <code>cdc.jar</code> , <code>foundation-rmi.jar</code> . <i>Note:</i> <code>lib/class-lib.jar</code> is only present for non-preloaded builds (<code>CVM_PRELOAD_LIB=false</code>).
<code>lib/libawtjpeg.so</code>	<i>PBP/PP only.</i> The Independent JPEG Group's JPEG codec.
<code>lib/libqtawt.so</code>	<i>PBP/PP only.</i> The AWT/Qt native bridge.

TABLE 3-1 Standard Files

File	Description
lib/content-types.properties	The MIME content type system property table used by the sun.net.www package. Each entry maps a MIME content type to a native application that can handle it. Files are associated with a MIME content type by either the MIME content type returned by an HTTP header or their file name extension.
lib/security/java.policy	System-wide security policies. ¹
lib/security/java.security	Master security properties. ¹
lib/zi/America/Los_Angeles	Time zone data files used by sun.util.calendar.ZoneInfoFile.
lib/zi/Asia/Calcutta	
lib/zi/Asia/Novosibirsk	
lib/zi/GMT	
lib/zi/ZoneInfoMappings	

¹ See *Inside Java 2 Platform Security, Second Edition: Architecture, API Design, and Implementation* by Li Gong (Addison-Wesley, 2003) for more information about Java SE security features.

3.2 CDC AMS Files

The CDC AMS version of the CDC Java runtime environment includes extra files for the mtask driver utility, the AMS and its application repository.

TABLE 3-2 CDC AMS Files

File	Description
bin/cvmc	The mtask driver utility controls JVM instances through the mtask command language described in Appendix B . See Appendix C for a complete description of the mtask driver utility.
lib/appmanager.jar	The AMS implementation is contained in two files that are usually loaded once in the first client JVM instance: <ul style="list-style-type: none">• a separate jar file contains the presentation mode• appmanager.jar contains the rest of the AMS implementation.
lib/appmanager-client.jar	Client JVM instances encapsulate managed applications. Each client JVM instance must load the com.sun.appmanager.client package in appmanager-client.jar so that it can be managed by CDC AMS. For convenience, this package is loaded once in the mtask server JVM instance and therefore each client JVM instance can access it through shared memory.
lib/ AwtPDA_PresentationMode.jar PBP_PresentationMode.jar	CDC AMS has a modular architecture that isolates the user experience functionality into a small set of classes called a presentation mode. CDC AMS includes two alternate presentation modes. PBP_PresentationMode.jar contains a very simple presentation mode based on Personal Basis Profile and AwtPDA_PresentationMode.jar contains a more featureful presentation mode based on Personal Profile and which includes an over-the-air (OTA) application deployment capability.
lib/j2me_xml_cdc.jar	A subset of the Java Web Services (JSR-172) RI that provides XML processing services for the OTA application provisioning feature in the Personal Profile presentation mode.

TABLE 3-2 CDC AMS Files (*Continued*)

File	Description
lib/security/ appmanager.security.constrained appmanager.security.permissive appmanager.security.policy	CDC AMS-specific security policy files.
repository/ apps/ icons/ menu/ preferences/ profiles/	The application repository is a mechanism for storing, updating and removing applications. In the modular architecture of CDC AMS, the application repository is an implementation of the com.sun.appmanager.apprepository.AppRepository interface. See Section 5.6, “Installing Applications” on page 5-10 for a description of the application repository.
/tmp/appmanager_cleanup_user /tmp/appmanager_server_user /tmp/appmanager_warmup_user /tmp/appmanager_app_user	These logfiles are created by CDC AMS to record information about running applications and server processes.

3.3 Security Files

TABLE 3-3 describes optional security files in versions of the CDC Java runtime environment that include the security optional packages. See *Inside Java 2 Platform Security: Architecture, API Design, and Implementation* by Li Gong (second edition, Addison-Wesley, 2003) for more information about Java SE security features.

TABLE 3-3 Security Files

File	Description
lib/jaas.jar	<i>Java Authentication and Authorization Service (JAAS) Optional Package</i> is a part of JSR-219 which is a framework for enforcing access control to resources using a CodeSource-based and Subject-based security model. jaas.jar contains the JAAS Optional Package implementation and the KeyStoreLoginModule authentication module, which is a subset of what is available in J2SE version 1.4.2.
lib/jce.jar lib/ext/sunjce_provider.jar lib/sunrsasign.jar	<i>Java Cryptography Extension (JCE) Optional Package</i> is a part of JSR-219 which extends the Java Cryptography Architecture (JCA) to include key generation and agreement, encryption and message authentication code (MAC) generation services. jce.jar contains the JCE Optional Package implementation which is fully compatible with J2SE version 1.4.2. sunjce_provider.jar contains the default ("SunJCE") provider implementation of the JCE service provider interface (SPI) and is fully compatible with J2SE version 1.4.2. Note that lib/ext is part of the extension class search path, but not part of the system class search path. See Section 4.3, "Class Search Path Basics" on page 4-4 for more information about class search paths. sunrsasign.jar contains the default ("SUN") provider implementation of the RSA signature SPI and is fully compatible with the SunJCE provider implementation in J2SE version 1.4.2. See "How to Implement a Provider for the Java Cryptography Architecture" in JSR-219.

TABLE 3-3 Security Files (Continued)

File	Description
lib/jsse-cdc.jar	<i>Java Secure Socket Extension (JSSE) Optional Package</i> is a part of JSR-219 which provides support for secure communication. jsse.jar contains both the JSSE Optional Package implementation and the default ("SunJSSE") provider implementation, which is fully compatible with the SunJSSE provider implementation in J2SE version 1.4.2.
lib/security/cacerts	Certificate authority (CA) keystore file. The default keystore password is "changeit". See keytool(1) for more information about how to use the Java SE SDK key and certificate management tool to change the keystore password.
lib/security/local_policy.jar	Security jurisdiction policy files.
lib/security/US_export_policy.jar	

3.4 Development Files

TABLE 3-4 describes files that can be used with developer tools like compilers and debuggers. These files are further described in [Chapter 8](#).

TABLE 3-4 Development Files

File	Description
lib/btclasses.zip	The CDC Java class library can be used for compiling application source code. <i>Note:</i> Because the contents of these archive files can vary depending on the selected build options, application development must be based on a target development version of the CDC Java class library. See the companion document <i>CDC Build System Guide</i> for information about how to build a target development version of the CDC Java class library.
lib/libdt_socket[_g].so lib/libjdwp[_g].so	The Java Debugger Wireline Protocol (JDWP) shared libraries are necessary for remote debugging.
hprof/	hprof profiler-related shared libraries, object files and build flags.

3.5

Test and Demonstration Files

TABLE 3-5 describes the test and demo programs. These are often included with the installation bundle, but are not necessary for operation.

TABLE 3-5 Test and Demonstration Files

File	Description
democlasses.jar	Demonstration applications that demonstrate profile-based functionality. This jar archive also contains the Java source code for these demo applications.
testclasses.zip	Test applications that can be used to quickly test the CDC Java runtime environment. The source code for these programs is located in <code>src/share/javavm/test</code> of the source code release. The simplest test programs to use are <code>HelloWorld</code> and <code>Test</code> .

Running Applications

The CDC Java runtime environment includes `cvm`, the CDC application launcher, for loading and executing Java applications. This chapter describes basic use of the `cvm` command to launch different kinds of Java applications, as well as more advanced topics like memory management and dynamic compiler policies.

4.1 Launching a Java Application

`cvm`, the CDC application launcher is similar to `java`, the Java SE application launcher. Many of `cvm`'s command-line options are borrowed from `java`. The basic method of launching a Java application is to specify the top-level application class containing the `main()` method on the `cvm` command-line. For example,

```
% cvm HelloWorld
```

By default, `cvm` looks for the top-level application class in the current directory. As an alternative, the `-cp` and `-classpath` command-line options can specify a list of locations where `cvm` can search for application classes. For example,

```
% cvm -cp /mylib:democlasses.jar HelloWorld
```

Here `cvm` searches for a top-level application class named `HelloWorld`, first in the directory `/mylib` and then in the archive file `democlasses.jar`. See [Section 4.3, “Class Search Path Basics” on page 4-4](#) for more information about class search paths.

The `-help` option displays a brief description of the available command-line options. [Appendix A](#) provides a complete description of the command-line options available for `cvm`.

4.2

Running Managed Applications (*Personal Basis Profile and Personal Profile only*)

Managed application models allow developers to offload the tasks of deployment and resource management to a separate application manager. The CDC Java runtime environment includes sample application managers for two different application models:

- The *applet application model* was one of the first success stories of Java technology. An applet contains dynamic web content that a user can view and manipulate through an applet viewer, typically built into a web browser.
- The *xlet application model* is similar in purpose to the applet application model, but different in design. The main differences between an xlet and an applet are that an xlet has a cleaner life cycle model and doesn't require an explicit dependency on AWT. These features make xlets more appropriate for embedded device scenarios like set-top boxes and PDAs.

4.2.1

Running an Applet (*Personal Profile only*)

The CDC Java runtime environment includes a simple applet launcher named `sun.applet.AppletViewer` which displays each applet in a separate frame. `AppletViewer` is a simplified version of the Java SE `appletviewer` utility (see <http://java.sun.com/j2se/1.4.2/docs/tooldocs/solaris/appletviewer.html>). The CDC version does not have a separate command-line utility and does not support all of the options available in the Java SE version, such as `-debug`, `-encoding`, and `-J`.

The basic command syntax to launch an applet with `AppletViewer` is:

```
% cvm sun.applet.AppletViewer URL
```

`URL` identifies an HTML file containing an APPLET, OBJECT or EMBED tag that identifies an applet. See <http://java.sun.com/j2se/1.4.2/docs/tooldocs/appletviewertags.html> for a description of the HTML tags supported by `AppletViewer`.

Here is a simple example of how to launch an applet based on the `DemoApplet` example in `democlasses.jar`:

```
% cvm sun.applet.AppletViewer personal/DemoApplet.html
```

4.2.2

Running an Xlet (*Personal Basis Profile and Personal Profile only*)

The CDC Java runtime environment includes a simple xlet manager named `com.sun.xlet.XletRunner`. Xlets can be graphical, in which case the xlet manager displays each xlet in its own frame, or they can be non-graphical. The basic command syntax to launch `XletRunner` is:

```
% cvm com.sun.xlet.XletRunner { \
    -name xletName \
    (-path xletPath | -codebase urlPath) \
    -args arg1 arg2 arg3 ...} \
    ...
% cvm com.sun.xlet.XletRunner -filename optionsFile
% cvm com.sun.xlet.XletRunner -usage
```

TABLE 4-1 describes `XletRunner`'s command-line options:

TABLE 4-1 XletRunner Command-Line Options

Option	Description
<code>-name xletName</code>	<i>Required</i> . Identifies the top-level Java class that implements the <code>javax.microedition.xlet.Xlet</code> interface.
<code>-path xletPath</code>	<i>Required</i> (or substituted with the <code>-codebase</code> option described below). Specifies the location of the target xlet with a local pathname. The path can be absolute or relative to the current directory. If the xlet is in a <code>JAR</code> or Zip archive file, then use the archive file name. <i>Note</i> : The xlet <i>must not</i> be found in the system class path, especially when running more than one xlet, because xlets must be loaded by their own class loader.
<code>-codebase urlPath</code>	<i>Optional</i> . Specifies the location of the target xlet with a URL. The <code>-codebase</code> option can be substituted for <code>-path</code> to provide a URL-formatted path instead of a local pathname.
<code>-args arg1 [arg2] [arg3] ...</code>	<i>Optional</i> . Passes additional runtime arguments to the xlet. Multiple arguments are separated by spaces.
<code>-filename optionsFile</code>	<i>Optional</i> . Reads options from an ASCII file rather than from the command line. The <code>-filename</code> option must be the first option provided to <code>XletRunner</code> .
<code>-usage</code>	Display a usage string describing <code>XletRunner</code> 's command-line options.

Here are some command-line examples for launching xlets with `com.sun.xlet.XletRunner`:

- To run `DemoXlet` in `democlasses.jar`:

```
% cvm com.sun.xlet.XletRunner \
  -name basis.DemoXlet \
  -path democlasses.jar
```

- To run an xlet with multiple command-line arguments:

```
% cvm com.sun.xlet.XletRunner \
  -name MyXlet \
  -path . \
  -args top bottom right left
```

- To run more than one xlet, repeat the `XletRunner` options:

```
% cvm com.sun.xlet.XletRunner \
  -name ServerXlet -path ./server \
  -name ClientXlet -path ./client
```

- To run an xlet whose compiled code is at the URL

<http://java.sun.com/xlets/demo/MyXlet.class>:

```
% cvm com.sun.xlet.XletRunner \
  -name MyXlet \
  -codebase http://java.sun.com/xlets/demo
```

- To run an xlet in a jar file named `xlet.jar` with the arguments `colorMap` and `blue`, use the following command line:

```
% cvm com.sun.xlet.XletRunner \
  -name StockTickerXlet \
  -path xlet.jar \
  -args colorMap blue
```

- To run an xlet with command-line options in an argument file:

```
% cvm com.sun.xlet.XletRunner -filename myArgsFile
```

`myArgsFile` contains a text line with valid `XletRunner` options:

```
-name StockTickerXlet -path xlet.jar -args colorMap blue
```

4.3

Class Search Path Basics

The Java runtime environment uses various search paths to locate classes, resources and native objects at runtime. This section describes the two most important search paths: the Java class search path and the native method search path.

4.3.1 Java Class Search Path

Java applications are collections of Java classes and application resources that are built on one system and then potentially deployed on many different target platforms. Because the file systems on these target platforms can vary greatly from the development system, Java runtime environments use the *Java class search path* as a flexible mechanism for balancing the needs of platform-independence against the realities of different target platforms.

The Java class search path mechanism allows the Java virtual machine to locate and load classes from different locations that are defined at runtime on a target platform. For example, the same application could be organized in one way on a MacOS system and another on a Linux system. Preparing an application's classes for deployment on different target systems is part of the development process. Arranging them for a specific target system is part of the deployment process.

The Java class search path defines a list of locations that the Java virtual machine uses to find Java classes and application resources. A location can be either a file system directory or a jar or Zip archive file. Locations in the Java class search path are delimited by a platform-dependent path separator defined by the `path.separator` system property. The Linux default is the colon ":" character.

The Java SE documentation¹ describes three related Java class search paths:

- The *system* or *bootstrap classes* comprise the Java platform. The *system class search path* is a mechanism for locating these system classes. The default system search path is *JRE/lib*.
- The *extension classes* extend the Java platform with optional packages like the JDBC Optional Package. The *extension class search path* is a mechanism for locating these optional packages. *cvm* uses the `-Xbootclasspath` command-line option to statically specify an extension class search path at launch time and the `sun.boot.class.path` system property to dynamically specify an extension class search path. The CDC default extension class search path is *CVM/lib*, with the exception of some of the provider implementations for the security optional packages described in [TABLE 3-3](#) which are stored in *CVM/lib/ext*. The Java SE default extension class search path is *JRE/lib/ext*.
- The *user classes* are defined and implemented by developers to provide application functionality. The *user class search path* is a mechanism for locating these application classes. Java virtual machine implementations like the CDC Java runtime environment can provide different mechanisms for specifying an Java class search path. *cvm* uses the `-classpath` command-line option to statically specify an Java class search path at launch time and the `java.class.path`

1. See the tools documentation at

<http://java.sun.com/j2se/1.4.2/docs/tooldocs/tools.html> for a description of the J2SDK tools and how they use Java class search paths.

system property to dynamically specify an user class search path. The Java SE application launcher also uses the CLASSPATH environment variable, which is *not* supported by the CDC Java runtime environment.

4.3.2 Native Method Search Path

The CDC HotSpot Implementation virtual machine uses the Java Native Interface¹ (JNI) as its native method support framework. The JNI specification leaves the platform-level implementation of native methods up to the designers of a Java virtual machine implementation. For the Linux-based CDC Java runtime environment described in this runtime guide, a JNI native method is implemented as a Linux shared library that can be found in the native library search path defined by the `java.library.path` system property.

Note – The standard mechanism for specifying the native library search path is the `java.library.path` system property. However, the Linux dynamic linking loader may cause other shared libraries to be loaded implicitly. In this case, the directories in the `LD_LIBRARY_PATH` environment variable are searched *without* using the `java.library.path` system property. One example of this issue is the location of the Qt shared library. If the target Linux platform has one version of the Qt shared library in `/usr/lib` and the CDC Java runtime environment uses another version located elsewhere, this directory must be specified in the `LD_LIBRARY_PATH` environment variable.

Here is a simple example of how to build and use an application with a native method. The mechanism described below is very similar to the Java SE mechanism.

- 1. Compile a Java application containing a native method.**

```
% javac -bootclasspath lib/btclasses.zip HelloJNI.java
```

- 2. Generate the JNI stub file for the native method.**

```
% javah -bootclasspath lib/btclasses.zip HelloJNI
```

- 3. Compile the native method library.**

```
% gcc HelloJNI.c -shared -I${CDC_SRC}/src/share/javavm/export \
-I${CDC_SRC}/src/linux/javavm/include -o libHelloJNI.so
```

This step requires the CDC-based JNI header files in the CDC source release.

- 4. Relocate the native method library in the test directory.**

```
% mkdir test
% mv libHelloJNI.so test
```

1. See the *Java Native Interface: Programmer's Guide and Specification* described in “[Related Documentation](#)” on page [xvi](#).

5. Launch the application.

```
% cvm -Djava.library.path=test HelloJNI
```

If the native method implementation is not found in the native method search path, the CDC Java runtime environment throws an `UnsatisfiedLinkError`.

4.4

Memory Management

The CDC Java runtime environment uses memory to operate the Java virtual machine and to create, store and use objects and resources. This section provides an overview of how memory is used by the Java virtual machine. Of course, the actual memory requirements of a specific Java application running on a specific Java runtime environment hosted on a specific target platform can only be determined by application profiling. But this section will provide useful guidelines.

4.4.1

The Java Heap

When it launches, the CDC Java runtime environment uses the native platform's memory allocation mechanism to allocate memory for native objects and reserve a pool of memory, called the *Java heap*, for Java objects and resources. The size of the Java heap is specified by the `-Xmxsize` and `-Xmssize` command-line options described in [TABLE A-1](#).

For example,

```
% cvm -Xmx7M MyApp
```

launches the application `MyApp` and sets the Java heap size to 7 MB.

- If the requested Java heap size is larger than the available memory on the device, the Java runtime environment exits with an error message:

```
% java -Xmx23000M MyApp
Invalid maximum heap size: -Xmx23000M
Could not create the Java virtual machine.
```

- If there isn't enough memory to create a Java heap of the requested size, the Java runtime environment exits with an error message:

```
% java -Xmx2300M MyApp
Error occurred during initialization of VM
Could not reserve enough space for object heap
```

- If the application launches and later needs more memory than is available in the Java heap, the CDC Java runtime environment throws an `OutOfMemoryError`.

- In the current implementation, the Java heap is not growable. So if only one of these command-line options is specified, its value will be used to set the Java heap size. This is compatible with previous versions of CDC where `-Xms` was the only available option for specifying the Java heap size. If both command-line options are specified and `-Xmx` is less than `-Xms`, the `-Xms` value is used for the maximum Java heap size and the `-Xmx` value will be ignored. Otherwise, `-Xmx` is used as the Java heap size. Note that future implementations may not adhere to these rules for determining Java heap size.

4.4.2 Garbage Collection

When a Java application creates an object, the Java runtime environment allocates memory out of the Java heap. And when the object is no longer needed, the memory should be recycled for later use by other objects and resources. Conventional application platforms require a developer to track memory usage. Java technology uses an automatic memory management system that transfers the burden of managing memory from the developer to the Java runtime environment.

So the Java runtime environment detects when an object or resource is no longer being used by a Java application, labels it as “garbage” and later recycles its memory for other objects and resources. This *garbage collection* (GC) system frees the developer from the responsibility of manually allocating and freeing memory, which is a major source of bugs with conventional application platforms.

GC has some additional costs, including runtime overhead and memory footprint overhead. However, these costs are small in comparison to the benefits of application reliability and developer productivity.

4.4.2.1 Garbage Collection in the CDC HotSpot Implementation

The Java Virtual Machine Specification does not specify any particular GC behavior and early Java virtual machine implementations used simple and slow GC algorithms. More recent implementations like the Java HotSpot Implementation virtual machine provide GC algorithms tuned to the needs of desktop and server Java applications. And now the CDC HotSpot Implementation includes a GC framework that has been optimized for the needs of connected devices.

The major features of the GC framework in the CDC HotSpot Implementation are:

- *Exactness*. Exact GC is based on the ability to track all pointers to objects in the Java heap. Doing so removes the need for object handles, reduces object overhead, increases the completeness of object compaction and improves reliability and performance.

- *Default Generational Collector.* The CDC HotSpot Implementation Java virtual machine includes a generational collector that supports most application scenarios, including the following:
 - general-purpose
 - excellent performance
 - robustness
 - reduced GC pause time
 - reduced total time spent in GC
- *Pluggability.* While the default generational collector serves as a general-purpose garbage collector, the GC plug-in interface allows support for device-specific needs. Runtime developers can use the GC plug-in interface to add new garbage collectors *at build-time* without modifying the internals of the Java virtual machine. In addition, starter garbage collector plug-ins are available from *Java Partner Engineering* (www.sun.com/software/jpe).

Note – Needing an alternate GC plug-in is rare. If an application has an object allocation and longevity profile that differs significantly from typical applications (to the extent that the application profile cannot be catered to by setting the GC arguments), and this difference turns out to be a performance bottleneck for the application, then alternate GC implementation may be appropriate.

4.4.2.2 Default Generational Collector

The default generational collector manages memory in the Java heap. **FIGURE 4-1** shows how the Java heap is organized into two heap generations, a *young* generation and a *tenured* generation. The generational collector is really a hybrid collector in that each generation has its own collector. This is based on the observation that most

Java objects are short-lived. The generational collector is designed to collect these short-lived objects as rapidly as possible while promoting more stable objects to the tenured generation where objects are collected less frequently.

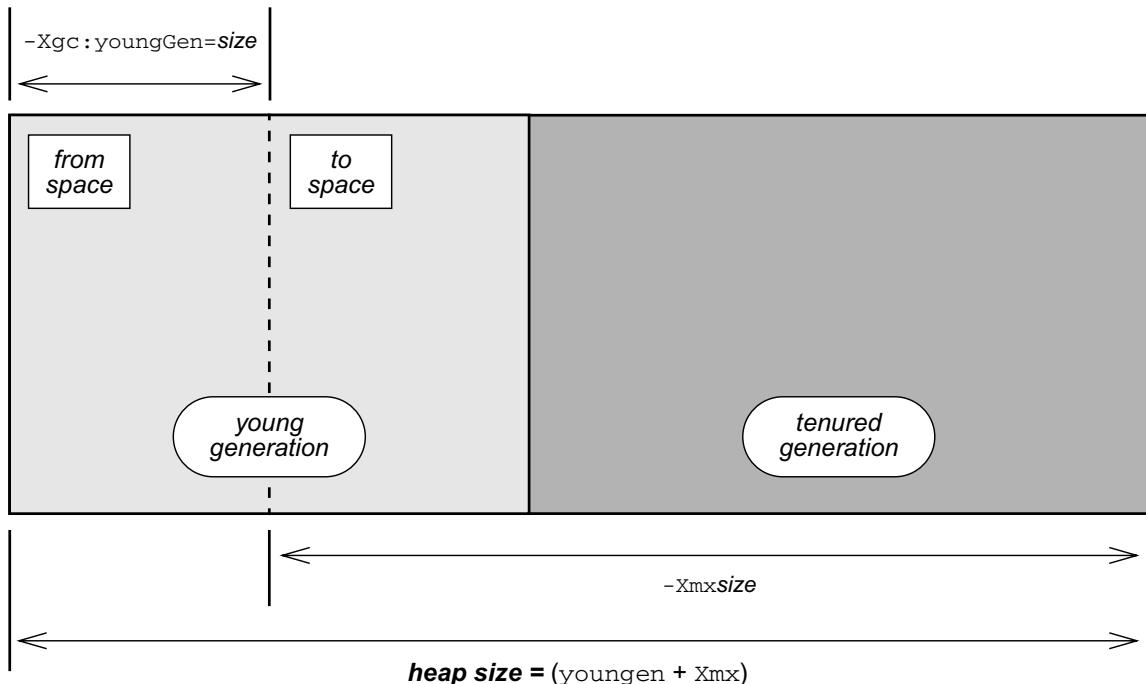


FIGURE 4-1 GC Generations

The young generation is based on a technique called *copying semispace*. The young generation is divided into two equivalent memory pools, the *from-space* and the *to-space*. Initially, objects are allocated out of the from-space. When the from-space becomes full, the system pauses and the young generation begins a collection cycle where only the live objects in the from-space are copied to the to-space. The two memory pools then reverse roles and objects are allocated from the “new” from-space. Only surviving objects are copied. If they survive a certain number of collection cycles (the default is 2), then they are promoted to the tenured generation.

The benefit of the copying semispace technique is that copying live objects across semispaces is faster than relocating them within the same semispace. This requires more memory, so there is a trade-off between the size of the young generation and GC performance.

The tenured generation is based on a technique called *mark compact*. The tenured generation contains all objects that have survived several copying cycles in the young generation. When the tenured generation reaches a certain threshold, the system pauses and it begins a full collection cycle where both generations go

through a collection cycle. The young generation goes through the stages outlined above. Objects in the tenured generation are scanned from their “roots” and determined to be live or dead. Next, the memory for dead objects is released and the tenured generation goes through a compacting phase where objects are relocated within the tenured generation.

The default generational garbage collector reduces performance overhead and helps collect short-lived objects rapidly, which increases heap efficiency.

4.4.2.3

Tuning Options

Two runtime options are available to control GC performance. These are described in [TABLE 4-2](#).

TABLE 4-2 GC Runtime Options

Option	Description
<code>-Xmxsize</code>	Maximum size of the Java heap
<code>-Xgc:youngGen=size</code>	Size of the young generation

The relative sizes of these generations can affect GC performance.

- `youngGen` should not be too small. If it is too small, partial GCs may happen too frequently. This causes unnecessary pauses and retain more objects in the tenured generation than is necessary because they don’t have time to age and die out between GC cycles.

The default size of `youngGen` is about 1/8 of the overall Java heap size.

- `youngGen` should not be too large. If it is too large, even partial GCs may result in lengthy pauses because of the number of live objects to be copied between semispaces or generations will be larger.

By default, the CDC Java runtime environment caps `youngGen` size to 1 MB unless it is explicitly specified on the command line.

- The total heap size needs to be large enough to cater for the needs of the application. This is very application-dependent and can only be estimated.

4.4.3

Class Preloading

The CDC HotSpot Implementation virtual machine includes a mechanism called *class preloading* that streamlines VM launch and reduces runtime memory requirements. The CDC build system includes a special build tool called `JavaCodeCompact` that performs many of the steps at build-time that the VM would normally perform at runtime. This saves runtime overhead because class

loading is done only once at build-time instead of multiple times at runtime. And because the resulting class data can be stored in a format that allows the VM to execute in place from a read-only file system (for example, Flash memory), it saves memory.

Note – It's important to understand that decisions about class preloading are generally made at build-time. See the companion document *CDC Build Guide* for information about how to use `JavaCodeCompact` to include Java class files with the list of files preloaded by `JavaCodeCompact` with the CDC Java runtime environment's executable image.

Class Preloading and Verification

Java class verification is usually performed at class loading time to insure that a class is well-behaved. This has both performance and security benefits. This section describes a performance optimization that avoids the overhead of Java class verification for some application classes.

One way to avoid the overhead of Java class verification is to turn it off completely:

```
% cvm -Xverify:none -cp MyApp.jar MyApp
```

This approach has the benefit of more quickly loading the application's classes. But it also turns off important security mechanisms that may be needed by applications that perform remote class loading.

Another approach is based on using `JavaCodeCompact` to *preload* an application's Java classes at build time. The application's classes load faster at runtime and other classes can be loaded remotely with the security benefits of class verification.

Note – `JavaCodeCompact` assumes the classes it processes are valid and secure. Other means of determining class integrity should be used at build-time.

The companion document *CDC Build Guide* describes how to use `JavaCodeCompact` to preload an application's classes so that they are included with the CDC Java runtime environment's binary executable image. Once built, the mechanism for running a preloaded application is very simple. Just identify the application without using `-cp` to specify the user Java class search path.

```
% cvm -Xverify:remote MyApp
```

The `remote` option indicates that preloaded and system classes will not be verified. Because this is the default value for the `-Xverify` option, it can be safely omitted. It is shown here to fully describe the process of preloading an application's classes.

4.4.4

Setting the Maximum Working Memory for the Dynamic Compiler

The `-Xjit:maxWorkingMemorySize` command-line option sets the maximum working memory size for the dynamic compiler. Note that the 512 KB default can be misleading. Under most circumstances the working memory for the dynamic compiler is substantially less and is furthermore temporary. For example, when a method is identified for compiling, the dynamic compiler allocates a temporary amount of working memory that is proportional to the size of the target method. After compiling and storing the method in the code buffer, the dynamic compiler releases this temporary working memory.

The average method needs less than 30 KB but large methods with lots of inlining can require much more. However since 95% of all methods use 30 KB or less, this is rarely an issue. Setting the maximum working memory size to a lower threshold should not adversely affect performance for the majority of applications.

4.5

Tuning Dynamic Compiler Performance

This section shows how to use `cvm` command-line options that control the behavior of the CDC HotSpot Implementation Java virtual machine's dynamic compiler for different purposes:

- Optimizing a specific application's performance.
- Configuring the dynamic compiler's performance for a target device.
- Exercising runtime behavior to aid the porting process.

Using these options effectively requires an understanding of how a dynamic compiler operates and the kind of situations it can exploit. During its operation the CDC HotSpot Implementation virtual machine instruments the code it executes to look for popular methods. Improving the performance of these popular methods accelerates overall application performance.

The following subsections describe how the dynamic compiler operates and provides some examples of performance tuning. For a complete description of the dynamic compiler-specific command-line options, see [Appendix A](#).

4.5.1 Dynamic Compiler Overview

The CDC HotSpot Implementation virtual machine offers two mechanisms for method execution: the *interpreter* and the *dynamic compiler*. The interpreter is a straightforward mechanism for executing a method's bytecodes. For each bytecode, the interpreter looks in a table for the equivalent native instructions, executes them and advances to the next bytecode. Shown in [FIGURE 4-2](#), this technique is predictable and compact, yet slow.

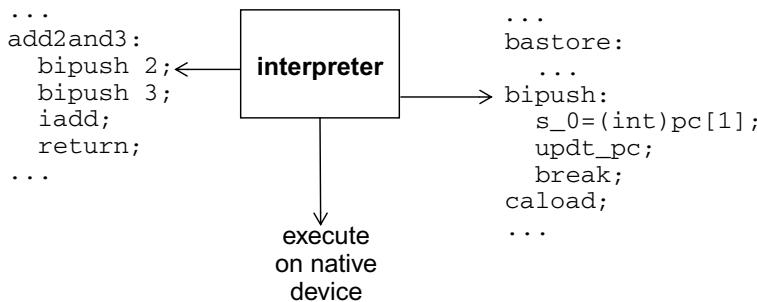


FIGURE 4-2 Interpreter-Based Method Execution

The dynamic compiler is an alternate mechanism that offers significantly faster runtime execution. Because the compiler operates on a larger block of instructions, it can use more aggressive optimizations and the resulting compiled methods run much faster than the bytecode-at-a-time technique used by the interpreter. This process occurs in two stages. First, the dynamic compiler takes the entire method's bytecodes, compiles them as a group into native code and stores the resulting native code in an area of memory called the *code cache* as shown in [FIGURE 4-3](#).

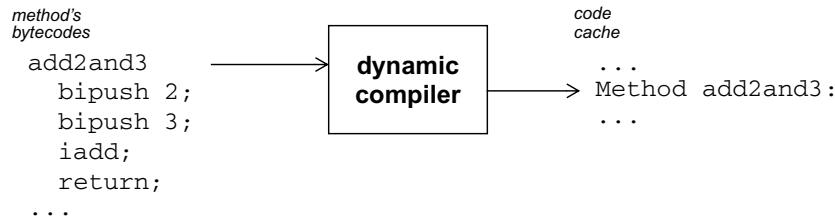


FIGURE 4-3 Compiling a Method

Then the next time the method is called, the runtime system executes the compiled method's native instructions from the code cache as shown in [FIGURE 4-4](#).



FIGURE 4-4 Executing a Compiled Method

The dynamic compiler cannot compile every method because the overhead would be too great and the start-up time for launching an application would be too noticeable. Therefore, a mechanism is needed to determine which methods get compiled and for how long they remain in the code cache.

Because compiling every method is too expensive, the dynamic compiler identifies important methods that can benefit from compilation. The CDC HotSpot Implementation Java virtual machine has a runtime instrumentation system that measures statistics about methods as they are executed. `cvm` combines these statistics into a single popularity index for each method. When the popularity index for a given method reaches a certain threshold, the method is compiled and stored in the code cache.

- The runtime statistics kept by `cvm` can be used in different ways to handle various application scenarios. To do this, `cvm` exposes certain weighting factors as command-line options. By changing the weighting factors, `cvm` can change the way it performs in different application scenarios. A specific combination of these options express a *dynamic compiler policy* for a target application. An example of these options and their use is provided in [Section 4.5.2.1, “Managing the Popularity Threshold” on page 4-16](#).
- The dynamic compiler has options for specifying code quality based on various forms of inlining. These provide space-time trade-offs: aggressive inlining provides faster compiled methods, but consume more space in the code cache. An example of the inlining options is provided in [Section 4.5.2.2, “Managing Compiled Code Quality” on page 4-17](#).
- Compiled methods are not kept in the code cache indefinitely. If the code cache becomes full or nearly full, the dynamic compiler decompiles the method by releasing its memory and allowing the interpreter to execute the method. An example of how to manage the code cache is provided in [Section 4.5.2.3, “Managing the Code Cache” on page 4-17](#).

4.5.2 Dynamic Compiler Policies

The `cvm` application launcher has a group of command-line options that control how the dynamic compiler behaves. Taken together, these options form *dynamic compiler policies* that target application or device specific needs. The most common are space-

time trade-offs. For example, one policy might cause the dynamic compiler to compile early and often while another might set a higher threshold because memory is limited or the application is short-lived.

TABLE A-8 describes the dynamic compiler-specific command-line options and their defaults. These defaults provide the best overall performance based on experience with a large set of applications and benchmarks and should be useful for most application scenarios. They might not provide the best performance for a specific application or benchmark. Finding alternate values requires experimentation, a knowledge of the target application's runtime behavior and requirements as well as an understanding of the dynamic compiler's resource limitations and how it operates.

The following examples show how to experiment with these options to tune the dynamic compiler's performance.

4.5.2.1 Managing the Popularity Threshold

When the popularity index for a given method reaches a certain threshold, it becomes a candidate for compiling. cvm provides four command-line options that influence when a given method is compiled: the popularity threshold and three weighting factors that are combined into a single popularity index:

- `climit`, the popularity threshold. The default is 20000.
- `bcost`, the weight of a backwards branch. The default is 4.
- `icost`, the weight of an interpreted to interpreted method call. The default is 20.
- `mcost`, the weight of transitioning between a compiled method and an interpreted method and vice versa. The default is 50.

Each time a method is called, its popularity index is incremented by an amount based on the `icost` and `mcost` weighting factors. The default value for `climit` is 20000. By setting `climit` at different levels between 0 and 65535, you can find a popularity threshold that produces good results for a specific application.

The following example uses the `-Xjit:option` command-line option syntax to set an alternate `climit` value:

```
% cvm -Xjit:climit=10000 MyTest
```

Setting the popularity threshold lower than the default causes the dynamic compiler to more eagerly compile methods. Since this will usually cause the code cache to fill up faster than necessary, this approach is often combined with a larger code cache size to avoid compilation/decompilation thrashing.

4.5.2.2

Managing Compiled Code Quality

The dynamic compiler can choose to inline methods for providing better code quality and improving the speed of a compiled method. Usually this involves a space-time trade-off. Method inlining consumes more space in the code cache but improves performance. For example, suppose a method to be compiled includes an instruction that invokes an accessor method returning the value of a single variable.

```
public void popularMethod() {  
    ...  
    int i = getX();  
    ...  
}  
public int getX() {  
    return x;  
}
```

`getX()` has overhead like creating a stack frame. By copying the method's instructions directly into the calling method's instruction stream, the dynamic compiler can avoid that overhead.

`cvm` has several options for controlling method inlining, including the following:

- `maxInliningCodeLength` sets a limit on the bytecode size of methods to inline. This value is used as a threshold that proportionally decreases with the depth of inlining. Therefore, shorter methods are inlined at deeper depths. In addition, if the inlined method is less than `value/2`, the dynamic compiler allows unquick opcodes in the inlined method.
- `minInliningCodeLength` sets the floor value for `maxInliningCodeLength` when its size is proportionally decreased at greater inlining depths..
- `maxInliningDepth` limits the number of levels that methods can be inlined.

For example, the following command-line specifies a larger maximum method size.

```
% cvm -Xjit:inline=all,maxInliningCodeLength=80 MyTest
```

4.5.2.3

Managing the Code Cache

On some systems, the benefits of compiled methods must be balanced against the limited memory available for the code cache. `cvm` offers several command-line options for managing code cache behavior. The most important is the size of the code cache, which is specified with the `codeCacheSize` option.

For example, the following command-line specifies a code cache that is half the default size.

```
% cvm -Xjit:codeCacheSize=256k MyTest
```

A smaller code cache causes the dynamic compiler to decompile methods more frequently. Therefore, you might also want to use a higher compilation threshold in combination with a lower code cache size.

The build option `CVM_TRACE_JIT=true` allows the dynamic compiler to generate a status report for when methods are compiled and decompiled. The command-line option `-Xjit:trace=status` enables this reporting, which can be useful for tuning the `codeCacheSize` option.

Application Management

CDC AMS includes two sample AMS implementations based on Personal Profile and Personal Basis Profile. These are collections of AMS services that are controlled by a user interface mechanism called a *presentation mode*. This chapter describes how to use these sample AMS implementations to demonstrate different aspects of application management.

- [Launching CDC AMS](#)
- [Exploring CDC AMS](#)
- [Launching Applications](#)
- [Switching Applications](#)
- [Terminating Applications](#)
- [Installing Applications](#)

5.1 Launching CDC AMS

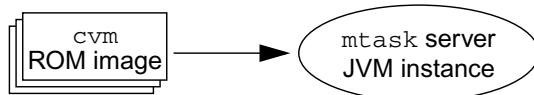
Note – See [Chapter 2](#), Installation for a description of how to install the CDC Java runtime environment on a target device. [Appendix G](#) provides platform-specific procedures for installing the CDC Java runtime environment on a Zaurus personal mobile tool. [Appendix H](#) provides platform-specific procedures for installing the CDC Java runtime environment on a Cobalt Qube.

The CDC AMS launch mechanism creates OS processes for the following:

- an mtask server JVM instance
- a client JVM instance that encapsulates the CDC AMS
- multiple client JVM instances that encapsulate managed applications

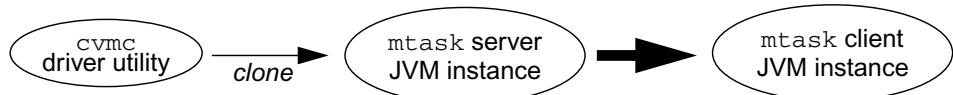
Launching CDC AMS is a three-stage procedure: First, we launch an `mtask` server JVM instance through the conventional JVM launch mechanism. This `mtask` server JVM instance remains in memory waiting for `mtask` commands. This procedure is shown in [FIGURE 5-1](#) where an `mtask` server JVM instance is created in memory through the conventional JVM launch mechanism based on loading and executing the CDC Java runtime environment from the `cvm` ROM image.

FIGURE 5-1 Launch an `mtask` Server JVM Instance



Next, we use the `cvmc` driver utility to send an `mtask` command to the server JVM instance to clone itself into a client JVM instance for running CDC AMS. [FIGURE 5-2](#) shows how this cloning procedure works. At the end of this cloning procedure, there are two JVM instances in memory, a server JVM instance which responds to `mtask` commands and a client JVM instance that encapsulates the AMS implementation.

FIGURE 5-2 `cvmc` Sends an `mtask` Command to Clone the Server JVM Instance



Later, CDC AMS can create more `mtask` client JVM instances for encapsulating individual applications, but that is managed interactively through the presentation mode and the AMS implementation. [FIGURE 5-3](#) illustrates how the AMS implementation can send an `mtask` command to the server JVM instance to clone itself. The AMS implementation could also send an `mtask` command to a client JVM instance to kill its process and end the application's lifecycle.

FIGURE 5-3 AppManager Sends an `mtask` Command to the Server JVM Instance



- [Appendix G](#) shows how to install CDC AMS on a Zaurus personal mobile tool. This includes installing the CDC Java runtime environment and integration with the Qtopia application environment. When this installation procedure is complete, you can launch CDC AMS by clicking on the `AwtPDA Appmanager` icon in the Java pane.
- [Appendix H](#) shows how to install CDC AMS on a Cobalt Qube.

5.2 Exploring CDC AMS

CDC AMS has two presentation modes with different user experiences. In general, the AwtPDA presentation mode is more featureful while the PBP presentation mode is more primitive. The procedures below are based on the Zaurus personal mobile tool running the AwtPDA presentation mode.

The Zaurus Qtopia application environment is organized like a conventional desktop environment. There are a set of icons organized into tabbed windows with a launch menu in the lower left-hand corner. For best results, the Zaurus Qtopia application environment should be in portrait display mode.

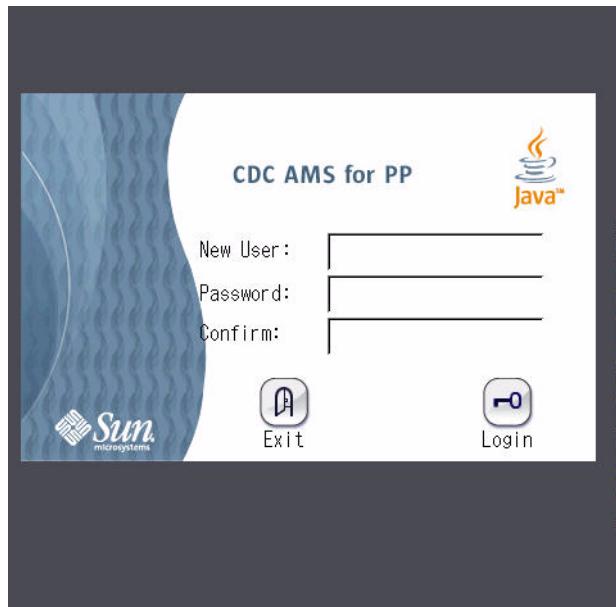
FIGURE 5-4 Zaurus Qtopia Application Environment



5.2.1 AwtPDA Presentation Mode (*Personal Profile only*)

After launching the AwtPDA presentation mode, the system displays a progress dialog followed by a login screen. The AwtPDA presentation mode includes a very simple user account system that supports a single user account per device. If there is currently no user account for the system, the user can create one and then login. Otherwise, the AwtPDA presentation mode displays a normal login screen.

FIGURE 5-5 AwtPDA Presentation Mode Login Screen

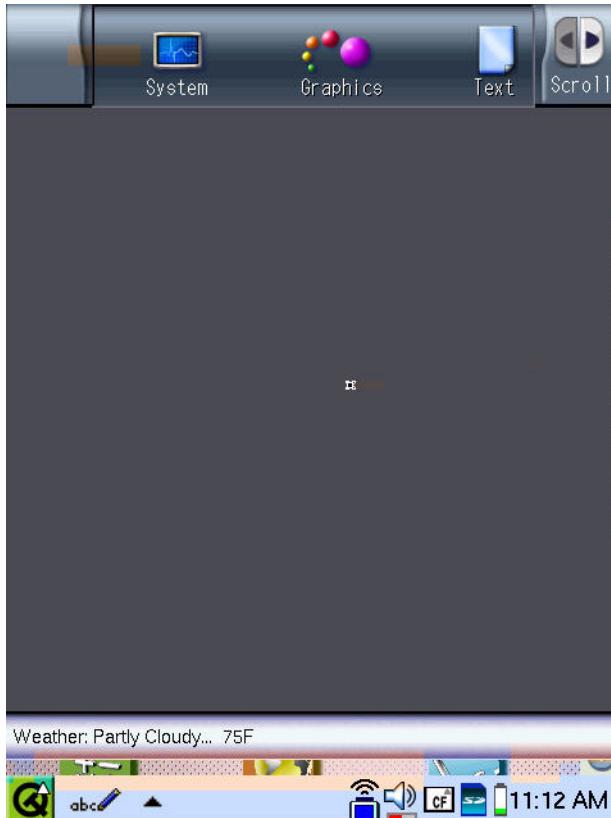


After logging in, the AWT presentation mode displays an application management screen which is organized into three areas:

- A scrollable *application strip* of icons on the top of the screen. The scroll buttons are on the right side of this strip. These scroll buttons control the display of a subset of the available application and system icons. A navigation button may be visible on the left side for navigating through a hierarchical icon folder.
- An *application pane* in the middle of the screen for displaying the GUI for an application.

- The *ticker area* at the bottom for displaying an ongoing newsfeed.

FIGURE 5-6 AwtPDA Presentation Mode Application Management Screen



The application strip organizes icons for applications and system utilities. A given icon can directly represent an application or utility or it can represent a folder containing other icons. To navigate between icons, use the scroll buttons on the right side of the application strip. After entering an icon group (like *System*), you can go up a level with the return button on the left side of the application strip.

The top-level icon groups are:

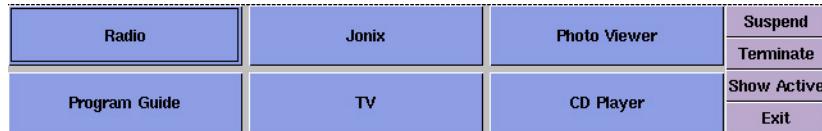
- *system* - exit, taskbar, application store, preferences
- *graphics* - graphics demo applications
- *text* - text applications
- *phone* - phone dialing application
- *tickers* - ongoing newsfeed application
- *games* - TicTacToe game
- *utils* - simple utilities like a clock and a helloworld xlet

5.2.2 PBP Presentation Mode (*Personal Basis Profile only*)

Note – The PBP presentation mode is a very simple AMS implementation that is designed around a simple button-based user-interface. It does not have much of the functionality of the AwtPDA presentation mode like the Taskbar utility or the OTA provisioning capability.

After launching, the PBP presentation mode displays the main application management screen, which is organized into a group of buttons on the bottom of the screen with six application buttons on the left and and four control buttons on the right. Pressing an application icon launches the corresponding application. Pressing a control button performs the corresponding system task.

FIGURE 5-7 PBP Presentation Mode Application Management Screen



There are six application buttons and four control buttons:

- Suspend - pauses the operation of an application
- Terminate - terminate an application
- Show Active - display a modal dialog with a list of active applications
- Exit - exit CDC AMS

5.3 Launching Applications

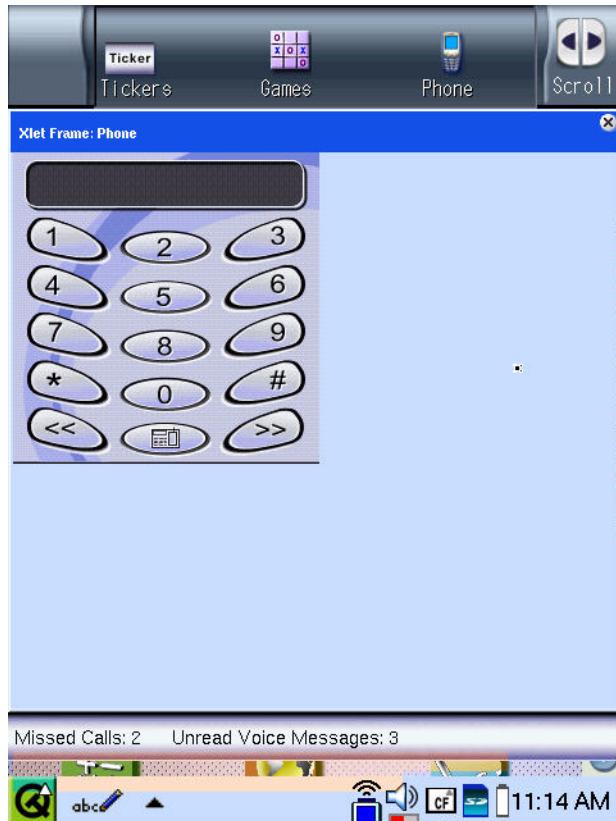
Launching an application is performed by clicking on the application's icon. The application will then display its GUI in the application pane until it is either terminated or another application is launched. In this latter case, the application will still be active and can be seen in the *System>Taskbar* (AwtPDA presentation mode) or *Show Active* (PBP presentation mode) system utilities.

5.3.1

AwtPDA Presentation Mode (*Personal Profile only*)

The phone xlet can be launched by first clicking on the right scroll button until the Phone icon appears and then clicking on the Phone icon itself. Several applications can be launched and run simultaneously but only one application is in the foreground and controls the application pane.

FIGURE 5-8 Phone Xlet



5.3.2 PBP Presentation Mode (*Personal Basis Profile only*)

The CD Player xlet can be launched by first clicking on the CDC Player button.

FIGURE 5-9 CD Player Xlet



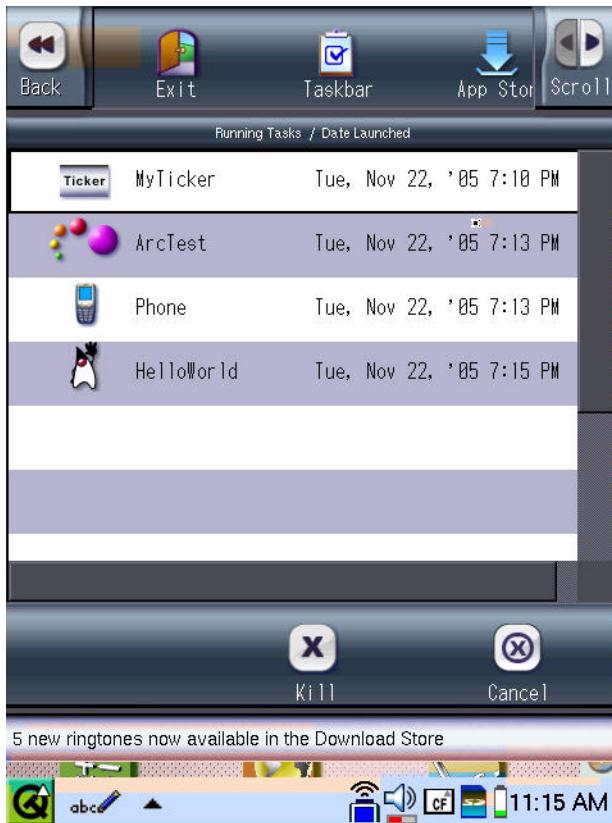
5.4

Switching Applications

The purpose of CDC AMS is to allow multiple applications to run concurrently. These applications share memory and CPU resources. However, only a single can be in the foreground and control the GUI panel. If one application is in the foreground and you want to interact with another, then you must switch the foreground application.

- In the AwtPDA presentation mode, an application can be selected through the System>Taskbar utility by selecting the application from the list and clicking the GoTo button. To select a different application, first launch the System>Taskbar utility. Then click on the target application to choose it. Then click on the Goto icon to select the application and bring it to the foreground.

FIGURE 5-10 Taskbar Utility



- In the PBP presentation mode, there is a very primitive application selection mechanism. Select the application by clicking on its button. This may cause an unlaunched application to launch.

5.5

Terminating Applications

The mechanics of terminating an application are similar to application switching. First, the application must be selected with the selection mechanism of the presentation mode and then the application can be terminated, again with the termination mechanism of the presentation mode.

- In the AwtPDA presentation mode, an application can be terminated through the System>Taskbar utility by selecting the application from the list and clicking the Kill button.
- In the PBP presentation mode, the application must be selected by clicking on its application button and terminated by clicking on the Terminate button.

5.6

Installing Applications

An application can be installed by adding its classes, icons and resources to the application repository. This can be done manually by copying files to the sub-directories in \$CVM/repository or automatically through the over-the-air application provisioning mechanism in the AwtPDA presentation mode. Once an application has been installed in the application repository, it will be visible in the CDC AMS presentation mode.

5.6.1 Manual Installation

The `$CVM/repository` directory has three sub-directories that contain application files. These are described in [TABLE 5-1](#).

TABLE 5-1 Application Repository Sub-Directories

Directory	Description
apps	Contains an subdirectory for each application. Within that directory is a <code>.jar</code> file that contains an application's classes and resources. For example, <code>repository/apps/Clock</code> contains an xlet for a clock application.
icons	Contains PNG-based icons for the application menu.
menu	Contains two kinds of files: <ul style="list-style-type: none">• a <code>.app</code> file is an application descriptor file with information used by the presentation manager to locate icons and other application resources.• a <code>.menu</code> file is a descriptor for sub-menus that describe application folders in the presentation mode. Both kinds of files use a simple key-value pair system to describe application location, icon location, etc.

Manually installing an application is simply a matter of installing an application's files in these three sub-directories.

5.6.2 OTA Provisioning

Note – The OTA provisioning feature of CDC AMS requires integration of the J2EE Client Provisioning RI server into the CDC build system and cannot be enabled with the binary implementation of CDC AMS.

[Appendix I](#) describes how to setup an OTA server based on using the J2EE Client Provisioning RI server to stage dynamic content like xlets. The CDC AMS client can interact with the OTA provisioning server to install applications into the application repository at runtime.

Security

Security is a principal feature of Java technology and an important requirement for mobile and enterprise applications. CDC includes the same security features that are in the Java SE platform. These include built-in security features of the Java programming language and virtual machine as well as a flexible security framework for more advanced application scenarios.

This chapter provides an overview of the security framework as well as an outline of the kinds of security procedures that might be performed at runtime. It is not meant to replace the security documentation available for the Java SE platform, but rather to supplement it and show how CDC and the JAAS, JCE and JSSE security optional packages are related to their counterparts in the Java SE platform.

TABLE 6-1 describes the security documentation for the Java SE platform.

TABLE 6-1 Security Documentation for the Java SE Platform

URL	Document	Description
http://java.sun.com/docs/books/security	<i>Inside Java 2 Platform Security</i>	Describes the Java security framework, including security architecture, deployment and customization. Chapter 12 describes deployment and runtime procedures.
http://java.sun.com/security	<i>Security and the Java Platform</i>	The main web page for Java security issues.
http://java.sun.com/docs/books/tutorial/security1.2	<i>Java Tutorial, Security Trail</i>	The <i>Java Tutorial</i> includes a security section that describes many of the security procedures for the Java platform. Because these are identical between CDC and the Java SE platform, they are not duplicated in this chapter.
http://java.sun.com/j2se/1.4.2/docs/guide/security	<i>Security</i>	Java SE platform security documentation.

6.1 Overview

The security framework shared by the Java SE platform and CDC is based on three key components:

- [Built-in Security Features](#)
- [Security Policy Framework](#)
- [Security Provider Architecture](#)

These provide a solid base for application and runtime security, a flexible mechanism for defining deployment-based security needs and a plug-in mechanism for supplying alternate security implementations.

6.1.1 Built-in Security Features

Java security is based on built-in language and VM security features that have been part of Java technology from its beginning:

- Strongly typed language (*runtime/compile-time/link-time*)
- Bytecode verification (*classloading-time*)
- Safety checks (*runtime*)
- Dynamic class loaders (*classloading-time*)

6.1.2 Security Policy Framework

A security policy controls how system resources are accessed by applications at runtime. The Java security framework includes both a default security policy and a mechanism for describing alternate security policies for application and deployment-specific needs. The main benefits of this security policy framework are:

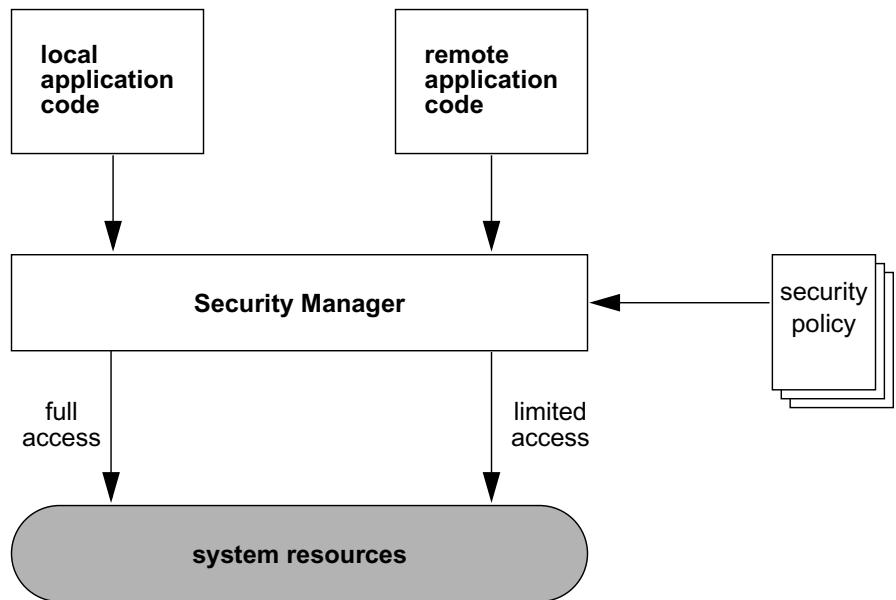
- Code-centric, not identity-centric architecture
- Security policies are described separately from both the applications they control and the Java runtime environment.
- Fine-grained access control at the package, class or field level
- Flexible permission mechanism
- Protection domains provide a layer of abstraction between permissions and code.

The main elements of a security policy are the following:

- `permission set`, a list of permissions granted to the code
- `codeBase`, the location from where the code is loaded
- `signedBy`, the author of the code
- `principal`, the identity of the entity running the code

FIGURE 6-1 illustrates the Java security model by showing how application code can be loaded from different sources: local and remote. The security manager controls access to system resources by comparing properties of the application code with the current security policy. The default security policy allows full access to local application code and limited access to remote application code. But other security policies are possible. For example, application code from a trusted yet remote source may be given greater access than untrusted code from a local source.

FIGURE 6-1 Java Security Policy Model



6.1.3 Security Provider Architecture

Beginning with version 1.2, the Java SE platform added some security optional packages that allow Java technology to adapt to more specific requirements of applications and deployments. These security optional packages include a security provider architecture that is *interoperable* because it is based on publicly available security standards, and *extensible* because alternate *security provider implementations* can be supplied without requiring modifications to application code.

For example, the JAAS, JCE and JSSE security optional packages include several *service provider interfaces* (SPIs) that describe the requirements of a security provider implementation. TABLE 3-3 describes the default Sun implementations for these security components.

6.2 Security Procedures

This section outlines the security procedures surrounding the Java security framework described in the previous section. Because these procedures are identical to the procedures used for the Java SE platform, this section just describes the procedure and indicates where to find the appropriate Java SE platform documentation.

6.2.1 Using Alternate Security Providers

From an administrator's perspective, the first step is to choose whether to install and use any alternate security providers. In most cases, the Sun default security providers described in [TABLE 3-3](#) are sufficient.

For a description of how to install alternate security providers, see *Inside Java 2 Platform Security, Second Edition*. Section 12.5, *Installing Provider Packages*, describes how to install alternate security providers.

6.2.2 Public Key Management

The JAAS optional package includes an extensible authentication framework that can use different forms of authentication. The default `LoginModule` is the `KeyStoreLoginModule`, which uses a protected database (Sun's JKS keystore file) to store public key data. Other forms of authentication are possible like smartcard or Kerberos.

The main tool for managing keystore files is `keytool(1)`, which is included in the Java SE platform toolset. `keytool` can be used for

- importing a key
- listing available keys
- replacing a key
- deleting a key

The default keystore file is in `lib/security/cacerts`, described in [TABLE 3-3](#).

For a description of how to use `keytool` to add and modify keystore entries, see Section 12.8, *Security Tools*, in *Inside Java 2 Platform Security, Second Edition*. The security trail in the *Java Tutorial* also covers how to use `keytool`.

6.2.3

Security Policy Management

Security policies are stored in security policy files. `policytool(1)` is a convenient GUI-based tool for managing security policies. With it, a system administrator can

- identify a keystore
- specify permissions
- specify a codebase

The location of the default security policy file is `lib/security.policy`, described in [TABLE 3-3](#). Alternate locations can be defined with the `-Djava.security.policy` command-line option.

For a description of how to use the `policytool` to manage security policies, see Section 12.8, *Security Tools*, in *Inside Java 2 Platform Security, Second Edition*. The security trail in the *Java Tutorial* also covers how to use `keytool`.

6.2.4

Seed Generation for Random Number Generation

The CDC Java runtime environment uses a native platform-provided source as an entropy gathering device for seed generation indicated by the `securerandom.source` system property. The Linux default for this system property is `file:/dev/random`.

On some Linux systems, `/dev/random` can block if it hasn't generated sufficient entropy before a random seed is needed and this can cause applications using `java.security.SecureRandom` to hang while waiting for the entropy pool to fill. To avoid this hang problem, the CDC Java runtime environment has a fallback mechanism to read from the `/dev/urandom` device when it determines that there isn't enough entropy for `/dev/random` to work promptly.

Note that `/dev/urandom` is not generally considered strong enough to support applications like keypair generation. If the strongest possible seed generation is required, this fallback mechanism can be disabled by setting the `microedition.securerandom.nofallback` property to `true`. Doing so may run the risk of application hangs on certain devices where the entropy pool is subject to early exhaustion.

Localization

The CDC Java runtime environment can be localized to support different languages and cultures. The following sections provide CDC-specific information for localization procedures:

- [Setting Locale System Properties](#)
- [Timezone Information Files](#)
- [Font Management \(Personal Basis Profile and Personal Profile only\)](#)

7.1

Setting Locale System Properties

In the CDC Java runtime environment, the locale system properties described in [TABLE 7-1](#) are set before `cvm` can parse its command-line arguments. Thus, it is not possible to change the locale by specifying these system properties on the `cvm` command-line with the `-Dproperty=value` option.

TABLE 7-1 Locale System Properties

System Property	Description
<code>user.language</code>	Two-letter language name code based on ISO 639.
<code>user.region</code>	Two-letter region name code based on ISO 3166.
<code>file.encoding</code>	Default character-encoding name based on the IANA Charset MIB.

On Linux, these properties are extracted from the `LANG` locale environment variable using the format `language_region.encoding`. The `user.language` property is obtained from the `language` code. The `user.region` property is obtained from the `region`

code. The `file.encoding` property is obtained from the `encoding` suffix. For example, to change the locale behavior of `cvm` on Linux, simply change the `LANG` locale environment variable to set the locale system properties.

```
% setenv LANG en_US.ISO8859_1
```

Therefore,

```
user.language = en
user.region = US
file.encoding = ISO8859_1
```

7.2 Timezone Information Files

The `lib/zi` directory contains a small set of example timezone information files. Additional files can be generated and placed in this directory. See the `javadoc(1)` comments for the `sun.util.calendar.ZoneInfoFile` class for information about generating alternate timezone information files.

7.3 Font Management (*Personal Basis Profile and Personal Profile only*)

In the CDC Java runtime environment, font management is a subset of the functionality provided by Java SE technology and is described below in [TABLE 7-2](#).

TABLE 7-2 Font Management Comparison

Feature	Java SE	CDC
Default font mapping between Java logical fonts and platform logical fonts is specified at build-time.	yes	yes
Logical font mapping in <code>lib/font.properties</code> file.	yes	no
Bundled Lucida fonts in <code>lib/fonts</code> .	yes	no
Application-specific fonts in an application's <code>jar</code> file.	yes	no

The six logical fonts available to a Java application are described in [TABLE 7-3](#).

In practice, the only way to specify alternate fonts is to remap the platform logical fonts. [Appendix F](#) contains a procedure for installing TrueType fonts and mapping them to logical platform fonts that are used by the CDC Java runtime environment for the Java logical fonts described in [TABLE 7-3](#).

TABLE 7-3 Logical Font Names

Java Logical Font	Qt Logical Font	Example	Description
default	Sans Serif	Courier	The default font is used when no other font is specified or if an attempt to match a font fails.
dialog	Sans Serif	Lucida Sans	A font for displaying fixed information within a dialog box or form.
dialoginput	Courier	Lucida Sans Typewriter	A font that is used for text fields within dialog boxes and forms that represent user input.
monospaced	Courier	Lucida Sans Typewriter	A non-proportional font where each character has the same width. This simplifies string width calculations for dialog boxes and forms.
sanserif	Sans Serif	Helvetica	A streamlined font that is simpler to render on low-resolution devices like computer monitors and faxes.
serif	Serif	Times Roman	A font with short lines at the end of the main strokes of a character to ease visual character recognition.

Developer Tools

One of the principal goals of CDC is to leverage conventional Java SE developer tools for use with CDC applications and devices. This chapter shows how to integrate the CDC Java runtime environment with Java SE developer tools like `javac`, `jdb` and `hprof`.

8.1

Compiling With `javac`

Compiling Java source code is a separate process from execution. All that is needed is application source code, a Java compiler like `javac` and an appropriate Java class library to compile against. In this way, a developer can compile a Java application on a desktop system and later download it onto a target device for testing or deployment.

This chapter first reviews the API relationship between the CDC and Java SE platforms. Then it shows how `javac` compiles a Java class for the Java SE platform and how this process changes for CDC. Finally, it shows how to compile an example CDC program.

8.1.1

CDC and Java SE

It is possible to take unmodified application software that was compiled for the Java SE platform and run it on a CDC Java runtime environment because the CDC Java virtual machine can load and execute Java classes that are compliant with the class specification for the Java SE platform.

FIGURE 8-1 describes the API relationship between the CDC and Java SE platforms. The two platforms have much in common, including most of the core Java class library. Differences between the CDC and Java SE APIs can cause discrepancies at runtime. These differences are based on the need to remove or change certain classes for memory, functionality or performance reasons.

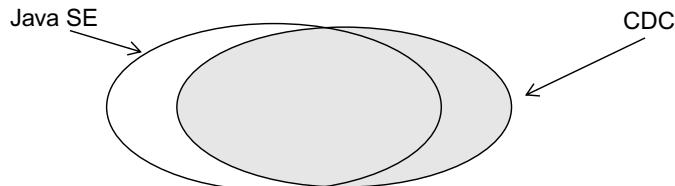


FIGURE 8-1 CDC and Java SE API Compatibility

There are four major differences between the CDC and Java SE platforms:

- Some Java SE packages, classes and methods have been removed because they are not appropriate for smaller devices. Compiling application source code against the Java SE class library may work, but the compiled classes may fail to run on a CDC Java runtime environment because the classes are not available at runtime.
- Some packages like `java.sql` are present in the Java SE platform but not in CDC, though they may be added as an optional package. In this case, compiling application source code against the Java SE class library may work but running the compiled classes against the CDC Java runtime environment may not.
- Most Java SE deprecated methods have been removed from CDC. For example, `java.awt.List.clear()` is deprecated in JDK version 1.1 and replaced with `java.awt.List.removeAll()`. In this case, compiling a Java SE application that uses this deprecated method against the CDC Java class library will cause `javac` to fail to compile because it cannot find the deprecated method.
- CDC includes CLDC compatibility classes that are not included in the Java SE class library. In this case, compiling CDC source code against the Java SE class library might cause `javac` to fail to compile because these compatibility classes are not present in the Java SE class library.

Therefore, in practice, it is best to recompile Java source code for a Java SE application against a CDC Java class library. Finally, the CDC Java class library is modular and can change based on the needs of a product design. Most of this modularity is based on profiles and optional packages. See [Section 1.6, “Java ME API Choices”](#) on page 1-6 for an explanation of how CDC APIs can vary.

8.1.2 Compiling Java Source Code for the Java SE Platform

[FIGURE 8-2](#) shows how `javac` compiles Java source code for the Java SE platform. When `javac` processes Java source code, it uses a Java class library to discover type information about the classes used in the source code. By default, this is the Java SE class library located in `jre/lib/rt.jar`.

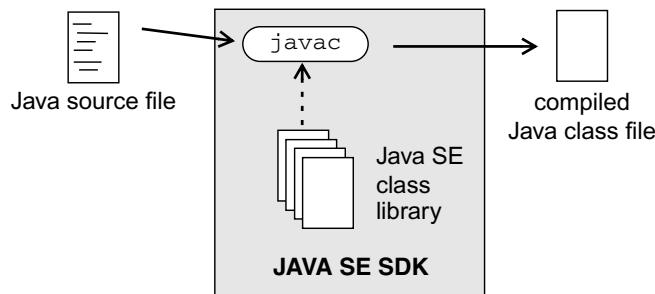


FIGURE 8-2 Compiling Java Source Code for the Java SE Platform

For example, when `javac` encounters a Java type reference like `java.util.BitSet`, it gets the type information from the Java SE class library at compile time. Later, at runtime, when the Java virtual machine creates an object of type `java.util.BitSet`, it also gets the type information from the Java SE class library.

8.1.3 Compiling Java Source Code for CDC

The same `javac` compiler used for developing Java SE applications can be used to compile Java source code for the CDC Java runtime system. The key is to use a different target Java class library to compile against. [FIGURE 8-3](#) shows how the `javac` compiler uses the `-bootclasspath` command-line option to specify an alternate target Java class library as a cross-compilation target.

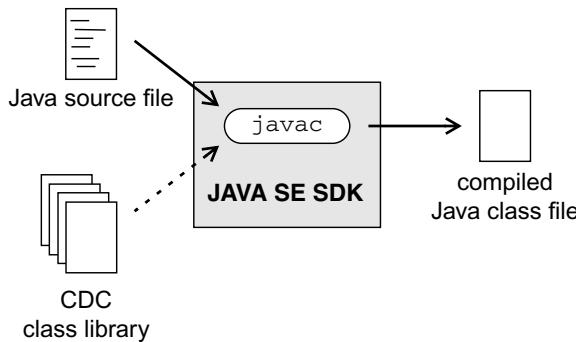


FIGURE 8-3 Compiling Java Source Code for CDC

The mechanics of using `javac` to compile Java source code for CDC differ slightly from those used for the Java SE platform.

8.1.4

Determining the Target Class Library

Section 1.5, “Java ME Technology Standards” on page 1-4, Section 1.6, “Java ME API Choices” on page 1-6, and FIGURE 1-2 show how the API functionality of a specific CDC product implementation can vary based on choices made at design time. Therefore, it is important to use a target development version of the CDC Java class library that represents the APIs available in the configuration, profile and optional packages on the target device.

Note – See the companion document *CDC Build System Guide* for information on how to build a target development version of the CDC Java class library which represents the combination of configuration, profile and optional packages for the target device.

8.1.5

Useful `javac` Command-Line Options

The J2SDK *Tools and Utilities* Web page (<http://java.sun.com/j2se/1.4.2/docs/tooldocs/tools.html>) describes the `javac` command-line options that control the cross-compilation process. These are described in the following subsections.

8.1.5.1 **`-classpath` *classpath***
Sets the user class search path, which is useful for compiling against third-party class libraries.

8.1.5.2 **`-bootclasspath` *classpath***
Sets the system class search path. With `javac`, this option overrides the Java SE class library and specifies an alternate target Java class library for cross-compilation like the target development version of the CDC Java class library.

8.1.5.3 **`-extdirs` *classpath***
Sets the extensions class search path for optional packages. The CDC default location is the `lib` directory, except for some security optional packages which are found in the `lib/ext` directory.

8.1.5.4 **`-source` *release***
Specifies the version of Java source code accepted. In practice, this controls the use of recently added Java programming language features. For example, J2SE 1.4 includes support for the `assert` keyword and J2SE 1.5 includes support for generics, which are not yet supported. The *release* argument can be set to 1.2, 1.3 or 1.4 for CDC application development.

8.1.5.5 **`-target` *version***
This option directs `javac` to generate Java class files for a specific version of the Java virtual machine. It is preferable to set the *version* value to 1.4, though values of 1.2 or 1.3 can also be used for CDC development.

8.1.5.6 **`-deprecation`**
Show a description of each use or override of a deprecated member or class. Without `-deprecation`, `javac` shows the names of source files that use or override deprecated members or classes.

8.1.6 Compiling an Example CDC Program

The example below demonstrates how to compile an application using the command-line option `-bootclasspath` argument to specify an alternate target Java class library:

```
% javac -target 1.4 -source 1.4 -bootclasspath \
/home/test-cdc/btclasses.zip MyApp.java
```

8.2 Debugging With `jdb`

A debugger like `jdb` can explore the relationships between the source code structure of an application, the behavior of its compiled code and the capabilities of the target Java runtime environment.

The CDC Java runtime environment supports debugging based on the Java Virtual Machine Debugger Interface (JVMDI) specification. This chapter describes the mechanics of attaching a remote debugger to a Java application running on a CDC Java runtime environment. The application can run on a target system while the debugger runs on a host development system connected over a network.

Note – Chapter 6 of the *CDC Build System Guide* describes how to build a version of the CDC Java runtime environment that supports Java debugging. Note that `cvm` must be compiled with both `CVM_JVMDI=true` and `CVM_JIT=false`.

8.2.1

Debug Command-Line Options

The debug version of `cvm` includes some extra command-line options described in **TABLE 8-1** that control debugging features. See <http://java.sun.com/j2se/1.4.2/docs/guide/jpda/conninv.html> for a complete list of `-Xrunjdwp` suboptions.

TABLE 8-1 `cvm` Debugging Options

Option	Description
<code>-Xdebug</code>	Run the VM in debugger mode.
<code>-Xrunjdwp</code> [<i>option1,option2...</i>]	Load the JDWP agent library (<code>libjdwp.so</code>). This library resides in the target VM and uses JVMDI and JNI to interact with it. It uses a transport and the JDWP protocol to communicate with a separate debugger application.
<code>transport=dt_socket</code>	Connect to the JVMDI debugger's front-end using a socket transport.
<code>server=y</code>	Start the VM in server mode and wait for the connection with a JVMDI debugger client.
<code>address=port</code>	Set the TCP port ID for the JDWP connection.

8.2.2

Running the Debug Version of `cvm`

Here's an example of how to launch a debug version of `cvm` on a remote target system for use with a host-based Java debugger. For this example, we assume the following:

- The target application is in `/net/MyApp`.
- The application is named `MyApplication`.
- The CDC Java runtime environment is correctly installed.
- A debug-capable version of `cvm` is in the shell's search path.

1. Remote login to the target system.

```
% ssh cdc-dev  
...
```

2. Change the current directory to the location of the target application.

```
% cd /net/MyApp
```

3. Launch `cvm` with the debug options.

```
% bin/cvm -Xrunjdwp:transport=dt_socket,server=y,address=8000 \
-Xdebug -Dsun.boot.library.path=jdwp/lib -cp democlasses.jar \
personal.DemoFrame
```

The `sun.boot.library.path` system property allows `cvm` to append to the shared library search path from the command line. This launches `cvm` in a server state where it waits for a connection with `jdb`, which is described in the next section.

8.2.3 Running `jdb` on the Host Development System

`jdb` is a JVMDI-based debugger that is included with the Java SE SDK. The example below shows how to attach `jdb` to an application running on a remote Java runtime environment.

This example assumes that J2SDK is properly installed and that `jdb` is in the shell's search path. Also, the source code for `MyApplication` should be in `/net/MyApp` so that `jdb` can access it.

- 1. Change the current directory to the location of the target application.**

```
% cd /net/MyApp
```

- 2. Launch `jdb` with the command-line options that identify the application on the target system.**

```
% jdb -attach cdc-dev:8000 -sourcepath src/personal/demo
```

`jdb` displays a command prompt.

- 3. Set a breakpoint.**

```
jdb> stop in MyApplication.main
```

- 4. Launch the application and let it run to the breakpoint.**

```
jdb> run
```

At this point, the application should be stopped at the first line of the top-level `main()` method.

- 5. Step through the application.**

```
jdb> step
```

See the `jdb` reference documentation

(<http://java.sun.com/products/jpda/doc>) for a list of options and commands or type `help` at the `jdb` command-line prompt.

8.3

Profiling with hprof

Profiling is the measurement of runtime data for a specific application on a target runtime system. Understanding the runtime behavior of an application allows the developer to identify performance-sensitive components when tuning an application's implementation or selecting runtime features.

The CDC HotSpot Implementation supports profiling based on the experimental Java Virtual Machine Profiler Interface (JVMPi) specification. Specifically, the JVMPi-based hprof profiling agent provides reports that include CPU usage, heap allocation statistics and monitor contention profiles.

Chapter 7 of the *CDC Build System Guide* describes how to build a version of the CDC Java runtime environment that supports profiling. This chapter describes the mechanics of using hprof to generate a simple profiling report.

Note – The JVMPi functionality in the CDC Java runtime environment is a subset of what is supported in the Java SE SDK. In particular, remote profiling is not supported.

8.3.1

Profiling Command-Line Options

The profiling version of cvm includes the `-Xrunhprof` command-line option described in [TABLE 8-2](#) that controls profiling features. See <http://java.sun.com/j2se/1.4.2/docs/guide/jvmpi/jvmpi.html> for more information.

TABLE 8-2 hprof Command-Line Options

Option	Default	Description
<code>-Xrunhprof[:help] [option=value, ...]</code>		Run the VM with hprof enabled
<code>heap=dump sites all</code>	<code>all</code>	Heap profiling
<code>cpu=samples times old</code>	<code>off</code>	CPU usage
<code>monitor=y n</code>	<code>n</code>	Monitor contention
<code>format=a b</code>	<code>a</code>	ASCII or binary output
<code>file=name</code>	<code>java.hpr</code> <code>of</code>	Write data to file <i>name</i> and append .txt for ASCII format

TABLE 8-2 hprof Command-Line Options

Option	Default	Description
<code>net=host:port</code>		Send data over a socket
<code>depth=size</code>	4	Stack trace depth
<code>cutoff=value</code>	0.0001	Output cutoff point
<code>lineno=y n</code>	y	Display line-numbers in traces
<code>thread=y n</code>	n	Thread in trace
<code>doe=y n</code>	y	Dump on exit

8.3.2 Running cvm With hprof

Here's an example of how to use hprof to profile an application.

```
% cvm -Xrunhprof:heap=all,cpu=samples,file=profile.txt MyApp
```

When the application terminates, the output file `profile.txt` contains the profile report.

cvm Reference

A.1 Synopsis

```
cvm [-options] class [options ...]  
cvm [-options] -jar jarfile [options ...]
```

A.2 Description

cvm launches a Java application. It does this by starting a Java virtual machine, loading its system classes, loading a specified application class, and then invoking that class's main method, which must have the following signature:

```
public static void main(String args[])
```

The first non-option argument to cvm is the name of the top-level application class with a fully-qualified class name that contains the main method. The Java virtual machine searches for the main application class, and other classes used, in three locations: the *system class path*, the *extension class path* and the *user class path*. See [Section 4.3, “Class Search Path Basics” on page 4-4](#) for more information about Java class paths. Non-option arguments after the main application class name are passed to the main method.

If the -jar *jarfile* command-line option is used, cvm launches the application in the jar file. The manifest of the jar file must contain a line of the form MainClass:*classname*. The *classname* string identifies the class having the main method which serves as the application's starting point.

[Section 4.1, “Launching a Java Application” on page 4-1](#) has more information about launching Java applications with cvm.

A.3 Options

cvm borrows some of its command-line options from java, the Java SE application launcher. Other options are unique to cvm and may require certain build options to enable the necessary runtime features.

Note – In most cases, the command-line options described below have identical behavior in standalone cvm mode and mtask mode. For command-line options with different behaviors in the two modes, the affected options are identified and the alternate behavior is explained below the standard behavior. In most of these cases, the client JVM instance *inherits* its behavior from the server JVM instance and the command-line options, if they have any effect at all, *override* the inherited behavior.

For command-line options that take a *size* parameter, the default units for size are bytes. Append the letter *k* or *K* to indicate kilobytes, or *m* or *M* to indicate megabytes.

TABLE A-1 describes the command-line options that are shared with the Java SE application launcher.

TABLE A-1 Java SE Command-Line Options

Option	Description
<code>-help</code>	Display usage information and exit. mtask: Only the client JVM instance exits.
<code>-showversion</code>	Display product version information and continue.
<code>-version</code>	Display product version information and exit. mtask: Only the client JVM instance exits.
<code>-Dproperty=value</code>	Set a system property value. See Appendix D for a description of security properties for CDC. mtask: Inherited. This command-line option overrides the behavior of the master JVM instance.
<code>-classpath classpath</code> <code>-cp classpath</code>	Specify an alternate user class path. ¹ The default user class path is the current directory. mtask: The inherited user class path for the master JVM instance is treated as read-only. <i>classpath</i> is appended to the default path. The <i>/a</i> and <i>/p</i> strings are ignored.
<code>-Xbootclasspath[/a /p]:classpath</code>	Specify the extension class path. ¹ <i>/a</i> appends <i>classpath</i> list to the default path. <i>/p</i> prepends <i>classpath</i> list to the default path. mtask: The extension class path inherited from the master JVM instance is treated as read-only. Therefore, <i>classpath</i> is appended to the default path whether or not the <i>/a</i> syntax is used.

TABLE A-1 Java SE Command-Line Options (*Continued*)

Option	Description
<code>-Xmssize</code>	<p>Set the minimum size of the memory allocation pool (heap). This value must be greater than 1000 bytes. The default value is 2M.</p> <p>NOTE: This option is ignored by the generational garbage collector, though it could be used by other garbage collectors.</p> <p><code>mtask</code>: Inherited. The heap size is fixed at launch time for the server JVM instance. So this command-line option is ignored by client JVM instances.</p>
<code>-Xmxsize</code>	<p>Set the maximum size of the memory allocation pool (heap). The default value is 5M.</p> <p><code>mtask</code>: Inherited. The heap size is fixed at launch time for the server JVM instance. So this command-line option is ignored by client JVM instances.</p>
<code>-Xsssize</code>	<p>Each Java thread has two stacks: one for Java code and one for native code. The maximum native stack size of the main thread is determined by the native application launcher (e.g. shell, OS, etc.). For subsequent threads, the maximum native stack size is set by the <code>-Xss</code> option, although this can be ignored by the underlying OS. See TABLE A-4 for a description of the command-line options for controlling the size of the Java stack.</p> <p>The default value is 0 which indicates that the value is actually set by the native environment.</p>

TABLE A-1 Java SE Command-Line Options (*Continued*)

Option	Description
<code>-enableassertions [:<package>... :<class>]</code> <code>-ea [:<package>... :<class>]</code>	<p>Enable Java assertions. These are disabled by default. With no arguments, this switch enables assertions for all user classes. With one argument ending in ..., the switch enables assertions in the specified package and any subpackages. If the argument is simply ..., the switch enables assertions in the unnamed package in the current working directory. With one argument not ending in ..., the switch enables assertions in the specified class.</p> <p>If a single command line contains multiple instances of these switches, they are processed in order before loading any classes. So, for example, to run a program with assertions enabled only in the package <code>com.wombat.fruitbat</code> (and any subpackages), the following command could be used:</p> <pre>% cvm -ea:com.wombat.fruitbat ... <MainClass></pre> <p>The <code>-enableassertions</code> and <code>-ea</code> switches apply to all class loaders and to system classes (which do not have a class loader). There is one exception to this rule: in their no-argument form, the switches do not apply to system. This makes it easy to turn on assertions in all classes except for system classes. The <code>-enablesystemassertions</code> option enables asserts in all system classes (that is, it sets the default assertion status for system classes to true). To run a program with assertions enabled in the package <code>com.wombat.fruitbat</code> but disabled in class <code>com.wombat.fruitbat.Brickbat</code>, the following command could be used:</p> <pre>% cvm -ea:com.wombat.fruitbat... \ -da:com.wombat.fruitbat.Brickbat <MainClass></pre> <p>mtask: For client JVM instances, each class has an assertion state that is inherited from the master JVM instance. Subsequently, within the client JVM instance, each new class gets its own assertion state.</p>

TABLE A-1 Java SE Command-Line Options (Continued)

Option	Description
<code>-disableassertions [:<package>... :<class>]</code> <code>-da [:<package>... :<class>]</code>	Disable Java assertions. This is the default behavior. With no arguments, <code>-disableassertions</code> or <code>-da</code> disables assertions. With one argument ending in ..., the option disables assertions in the specified package and any subpackages. If the argument is simply ..., the switch disables assertions in the unnamed package in the current working directory. With one argument not ending in ..., the switch disables assertions in the specified class. The <code>-disableassertions</code> and <code>-da</code> switches apply to all class loaders and to system classes that do not have a class loader. There is one exception to this rule: in their no-argument form, the switches do not apply to system. This makes it easy to turn on assertions in all classes except for system classes. A separate switch is provided to enable assertions in all system classes. See the description of the <code>-enablesystemassertions</code> option.
<code>-enablesystemassertions</code> <code>-esa</code>	Enable assertions in all system classes (sets the default assertion status for system classes to true). <code>mtask</code> : For client JVM instances, each class has an assertion state that is inherited from the master JVM instance. Subsequently, within the client JVM instance, each new class gets its own assertion state.
<code>-disablesystemassertions</code> <code>-dsa</code>	Disable assertions in all system classes.

1 See Section 4.3, “Class Search Path Basics” on page 44 and <http://java.sun.com/j2se/1.4.2/docs/tooldocs/tools.html> for more information about class search paths.

TABLE A-2 describes the CDC-specific command-line options.

TABLE A-2 CDC-Specific Command-Line Options

Option	Description
<code>-fullversion</code>	Display build version information and exit. <code>mtask</code> : Only the client JVM instance exits.
<code>-XbuildOptions</code>	Display build options and exit. <code>mtask</code> : Only the client JVM instance exits.
<code>-XshowBuildOptions</code>	Display build options and continue.
<code>-XappName=value</code>	Specify the application name for QPE. This is used to identify the <code>cvm</code> process for native application management and control.

TABLE A-2 CDC-Specific Command-Line Options (*Continued*)

Option	Description
<code>-Xverify:[all remote none]</code>	Perform class verification. <ul style="list-style-type: none">• <code>all</code> verify all classes.• <code>remote</code> verify all but preloaded and system classes.• <code>none</code> don't perform class verification. The default value is <code>remote</code> . If <code>-Xverify</code> is used without any arguments, the value is <code>all</code> . <code>mtask</code> : Inherited. This command-line option is only affects classes loaded in the client JVM instance.
<code>-XfullShutdown</code>	Make sure all resources are freed and the VM destroyed upon exit. This is the default for non-process-model operating systems, but is not needed for process-model operating systems, such as Linux. <code>mtask</code> : This command-line option is ignored by both server and client JVM instances.
<code>-Xgc:suboption</code>	Specify GC-specific options. The default GC is the generational garbage collector described in Chapter 4 . See TABLE A-3 for a description of the suboptions. Other garbage collectors are unsupported.
<code>-Xopt:suboption</code>	Control the Java stack. See TABLE A-4 for a description of the suboptions. The different suboptions can be appended into a single argument with name/value pair separated by commas.
<code>-Xserver:suboption</code>	Launch a server JVM instance (<code>mtask server</code>) to act as an instance factory for generating client JVM instances (<code>mtask clients</code>) through process cloning. Each <code>mtask</code> client inherits the <code>mtask</code> server's warmed up state. This feature supports the CDC AMS process-based application management infrastructure. See TABLE A-11 for a description of the suboptions. <code>mtask</code> : This command-line option is ignored by <code>mtask</code> clients.
<code>-Xtrace:flags</code>	Turn on trace flags. TABLE A-5 shows the hexadecimal values to turn on each trace flag. To turn on multiple flags, bitwise-OR the values of all the flags you wish to turn on, and use that result as the <code>-Xtrace</code> value. Requires the <code>CVM_TRACE=true</code> build option. (Unsupported.) <code>mtask</code> : This command-line option overrides the behavior inherited from the master JVM instance.

[TABLE A-3](#) describes the suboptions for the `-Xgc` command-line option.

TABLE A-3 `-Xgc : suboption`

Option	Description
<code>maxStackMapsMemorySize=size</code>	Set the size of the stack map cache. The default value is <code>0xFFFFFFFF</code> .
<code>stat</code>	Collect and display garbage collection statistics.
<code>youngGen=size</code>	Set the size of the young object generation. NOTE: this option is specific to the default generational collector. The default value is <code>1M</code> .
	<code>mtask</code> : Inherited. The <code>youngGen</code> size is fixed at launch time for the server JVM instance. So this command-line option is ignored by client JVM instances.

[TABLE A-4](#) describes the suboptions for the `-Xopt` command-line option, which controls the size of the Java stack. This option is useful for runtime development purposes only and is unsupported.

TABLE A-4 `-Xopt : suboption`

Suboption	Description
<code>stackMinSize=size</code>	Set the initial size of the Java stack, from <code><32...65536></code> . The default for JIT-based systems is <code>3K</code> and the default for non-JIT based systems is <code>1K</code> .
<code>stackMaxSize=size</code>	Set the maximum size of the stack, from <code><1024...1048576></code> . The default for <code>128K</code> .
<code>stackChunkSize=size</code>	Set the amount the stack grows when it needs to expand <code><32...65536></code> . The default for JIT-based systems is <code>2K</code> and the default for non-JIT based systems is <code>1K</code> .

[TABLE A-5](#) describes the flags used by the `-Xtrace` command-line option. This option is useful for runtime development purposes only and is unsupported.

TABLE A-5 `-Xtrace : flags (OI only, unsupported)`

Value	Description
<code>0x00000001</code>	Opcode execution.
<code>0x00000002</code>	Method execution.
<code>0x00000004</code>	Internal state of the interpreter loop on method calls and returns.
<code>0x00000008</code>	Fast common-case path of Java synchronization.

TABLE A-5 `-Xtrace:flags (OI only, unsupported) (Continued)`

Value	Description
0x00000010	Slow rare-case path of Java synchronization.
0x00000020	Mutex locking and unlocking operations.
0x00000040	Consistent state transitions. Garbage Collection (GC)-safety state only.
0x00000080	GC start and stop notifications.
0x00000100	GC root scans.
0x00000200	GC heap object scans.
0x00000400	GC object allocation.
0x00000800	GC algorithm internals.
0x00001000	Transitions between GC-safe and GC-unsafe states.
0x00002000	Class static initializers.
0x00004000	Java exception handling.
0x00008000	Heap initialization and destruction, global state initialization, and the safe exit feature.
0x00010000	Read and write barriers for GC.
0x00020000	Generation of GC maps for Java stacks.
0x00040000	Class loading.
0x00080000	Class lookup in VM-internal tables.
0x00100000	Type system operations.
0x00200000	Java code verifier operations.
0x00400000	Weak reference handling.
0x00800000	Class unloading.
0x01000000	Class linking.

[TABLE A-6](#) describes the command-line options available with the CVM_JVMDI build option. See [Chapter 8](#) for an example of how to use these command-line options.

TABLE A-6 JVMDI Options

Option	Description
<code>-Xdebug</code>	Enable VM-level debugging support.
<code>-Xrunjdwp: [option1,option2...]</code>	Load the JDWP library with the specified options. See TABLE 8-1 .

[TABLE A-7](#) describes the command-line options available with the CVM_JVMPI build option. See [Chapter 8](#) for an example of how to use these command-line options.

TABLE A-7 JVMPI Options

Option	Description
<code>-Xrunhprof:[help] [option=value, ...]</code>	Enable hprof profiling support.

[TABLE A-8](#) describes the command-line options available with the CVM_JIT=true build option. See [Chapter 4](#) for an example of how to use these command-line options.

TABLE A-8 `-Xjit:options (OI only)`

Option	Default	Description
<code>bcost=cost</code>	4	Cost of a backwards branch, between <0...32767>.
<code>climit=cost</code>	20000	The popularity threshold for a given method, between <0...65535>. The VM compares a per-method count based on <code>bcost</code> , <code>icost</code> and <code>mcost</code> against this threshold to determine when to compile a given method.
<code>codeCacheSize=value</code>	512k	Size of code cache where compiled methods are stored, between <0...32M>.
		<code>mtask</code> : Inherited. The cache size is fixed at launch time for the server JVM instance. So this command-line option is ignored by client JVM instances.
<code>compile=suboption</code>	policy	When to compile methods. See TABLE A-10 for descriptions of the suboptions for <code>compile</code> . The default policy is based on the suboption defaults listed in this table.
<code>icost=cost</code>	20	Cost of an interpreted-to-interpreted method call, between <0...32767>.
<code>inline=suboption</code>	all	Perform method inlining when compiling. See TABLE A-9 for descriptions of the suboptions for <code>inline</code> .
<code>lowerCodeCacheThreshold=percentage</code>	90%	Lower code cache threshold, between <0%...100%>. The dynamic compiler decompiles methods until the code cache reaches this threshold.
<code>maxCompiledMethodSize=value</code>	65535	Maximum size of a compiled method, between <0...64K>.

TABLE A-8 `-Xjit:options (OI only) (Continued)`

Option	Default	Description
<code>maxInliningCodeLength=value</code>	68	Maximum size of an inlined method, between <code><0...1000></code> . This value is used as a threshold that proportionally decreases with the depth of inlining. Therefore, shorter methods are inlined at deeper depths. In addition, if the inlined method is less than <code>value/2</code> , the dynamic compiler allows unquick opcodes in the inlined method.
<code>maxInliningDepth=value</code>	12	Maximum inlining depth of inlined methods/frames, between <code><0...1000></code> .
<code>maxWorkingMemorySize=value</code>	512k	Maximum working memory size for the dynamic compiler, between <code><0...64M></code> . See Section 4.4.4, "Setting the Maximum Working Memory for the Dynamic Compiler" on page 4-13 .
<code>mcost=cost</code>	50	Cost for transitioning between a compiled method and an interpreted method, and vice versa. Between <code><0..32767></code> .
<code>minInliningCodeLength=value</code>	16	The floor value for <code>maxInliningCodeLength</code> when its size is proportionally decreased at greater inlining depths.
<code>policyTriggeredDecompilations=boolean</code>	true	Policy triggered decompilations. If <code>false</code> , then never decompile a method to make room for more compilations. Methods remain compiled until the class is unloaded, even if the code cache is full.
<code>trace=suboption</code>		Set dynamic compiler trace options. See TABLE A-12 .
<code>upperCodeCacheThreshold=percentage</code>	95	Upper code cache threshold, between <code><0%...100%></code> . The dynamic compiler starts decompiling methods during a GC when the code cache passes this threshold unless <code>policyTriggeredDecompilations=false</code> .
<code>XregisterPhis=boolean</code>	true	<i>Unsupported.</i>
<code>XcompilingCausesClassLoader=boolean</code>	false	<i>Unsupported.</i>

[TABLE A-9](#) describes the command-line options for selecting when to inline methods.

TABLE A-9 `-Xjit:inline=suboption (OI only)`

Suboption	Description
all	Enable all the options listed below to perform inlining whenever possible. The default.
none	Do not perform inlining.
virtual	Perform inlining on virtual methods.
nonvirtual	Perform inlining on nonvirtual methods.
vhints	Virtual hints. Use hints gathered while interpreting a method to choose a target method to get inlined when an <code>invokevirtual</code> opcode is compiled.
ihints	Interface hints. Use hints gathered while interpreting a method to choose a target method for inlining when an <code>invokeinterface</code> opcode is compiled.
Xvsync	Inline virtual synchronized methods. Off by default. <i>Unsupported</i> .
Xnvsync	Inline non-virtual synchronized methods. Off by default. <i>Unsupported</i> .
Xdopriv	Inline privileged methods specified by <code>java.security.AccessController.doPrivileged()</code> . On by default. <i>Unsupported</i> .

[TABLE A-10](#) describes the top-level command-line options that control dynamic compiler policies.

TABLE A-10 `-Xjit:compile=suboption (OI only)`

Suboption	Description
policy	Compile according to existing compilation policy parameters such as <code>icost</code> and <code>climit</code> . The default.
all	Compile all methods aggressively. <i>Note:</i> this hurts performance and should be used only for testing the dynamic compiler.
none	Do not compile any methods.

[TABLE A-11](#) describes the suboptions for the server JVM instance.

TABLE A-11 `-Xserver:suboption (CDC AMS only)`

Suboption	Description
port	The TCP port that the server JVM instance is bound to.
initClasses	The class list file for preinitialization.
precompileMethods	The method list file for precompilation.

TABLE A-12 describes the command-line options for controlling dynamic compiler tracing. These options require a build with CVM_TRACE_JIT=true. These options are experimental and unsupported.

TABLE A-12 `-Xjit:trace=option` (*OI only*)

Suboption	Description
bctoir	Print information regarding the conversion of Java bytecodes to the JIT internal representation (IR), including a complete dump of all IR nodes.
codegen	Print the generated code in a format similar to the assembler language of the target processor. If the build option CVM_JIT_DEBUG=true, then this also prints the JavaCodeSelect rule used to generate the code interspersed with the generated code.
inlining	Print method inlining information during the bytecode to IR pass, such as which methods were inlined and which ones were not.
irop	Print information about optimizations done in the bytecode to IR pass.
osr	Print a message when compilation of a method is triggered by on stack replacement (OSR).
stats	Print statistics gathered during compilation.
status	Print a line of status each time a method is compiled. The output includes the name of the method and whether or not it was compiled successfully.

mtask Command Language Reference

This appendix describes the `mtask` command language in some detail. Both the `cvmc` driver utility and the CDC AMS implementation use the same `mtask` command language to launch, initialize and control JVM instances that encapsulate both the CDC AMS implementation and the applications it manages.

B.1 Launch Command Language

TABLE B-1 Launch Commands

Command	Description
BROADCAST <i>message</i>	Send the text string <i>message</i> to all mtask-based JVM instances.
CHILDREN_EXITED	Test whether any launched JVM instances have exited since the last time this command was executed.
JAPP	main() -based application-model specific launcher.
JDETACH <i>vmArgs className args</i>	Asynchronously launch a new client JVM instance that is detached from the server JVM instance (so that LIST, BROADCAST, KILLALL, etc. do not see this application). This is typically used during bootstrapping to launch the client JVM instance that encapsulates the appmanager and presentation mode implementations.
JEXIT	Terminate the mtask server JVM instance.
JSYNC <i>vmArgs className args</i>	Launch a client JVM instance and block the mtask server JVM instance until the client JVM instance task exits.
JXLET	Xlet application-model specific launcher.
KILL <i>taskID</i>	Terminate a task with the given <i>taskID</i> .
KILLALL	Terminate all launched tasks.
LIST	List currently running applications.
MESSAGE <i>taskID message</i>	Send the text string <i>message</i> to the mtask-based JVM instance identified by the process <i>taskID</i> .
S <i>vmArgs className args</i>	Execute <i>className</i> with <i>args</i> within the mtask server JVM instance. This is done in-place without using <code>fork()</code> and <code>execute()</code> . This command is typically used for warming up mtask server JVM instance.
SETENV <i>key=value</i>	Set the value of an environment variable in the mtask server JVM instance for subsequent use by launched applications.
TESTING_MODE <i>filePrefix</i>	For all subsequent launched apps, dump <i>filePrefix/stdout.taskId</i> , <i>filePrefix/stderr.taskId</i> and <i>filePrefix/exitcode.taskId</i> to correspond to the standard output, standard error and exit code for the task <i>taskId</i> . <i>filePrefix</i> is the path of the directory containing these logfiles.

B.2

Warming Up the mtask Server

mtask provides two ahead-of-time mechanisms for warming up the server JVM instance. Performing these warmup tasks once for the server JVM instance allows subsequent client JVM instances to benefit from these warmup mechanisms with no runtime cost by accessing the shared memory pages of the server JVM instance.

The two warmup mechanisms are pre-initializing classes and pre-compiling methods. The CDC Java runtime environment includes two files for configuring warmup in the `repository/profiles` directory. Each file contains a sample list of classes or methods for warmup.

- The *class preinitialization list* is located in `repository/profiles/classesList.txt`.
- The *class initialization list* is located in `repository/profiles/methodsList.txt`.

The syntax for these warmup list files is very simple. Each line represents a single class or method. Lines that begin with a hash mark (#) are commented out and ignored.

cvmc Reference

C.1 Synopsis

`cvmc [... cvmc options ...] [... cvm options ...]`

C.2 Description

`cvmc` is a driver utility for managing `mtask`-based JVM instances. Its main purpose is to perform tasks necessary to bootstrap a CDC AMS implementation. In addition, it can perform tasks that are performed internally by a CDC AMS implementation, such as application termination:

- warming up classes and methods
- cloning client JVM instances
- controlling the lifecycles of JVM instances
- monitoring JVM instances

`cvmc` communicates with `mtask`-based JVM instances through an established TCP port.

C.3 Options

`cvmc` recognizes two kinds of command-line options: `cvmc`-specific command-line options and `cvm` command-line options that are passed to a client JVM instance.

TABLE C-1 describes the cvmc command-line options.

TABLE C-1 cvmc Command-Line Options

Option	Description
<code>-childrenexited</code>	Equivalent to the CHILDREN_EXITED mtask command described in TABLE B-1.
<code>-command command</code>	Send <i>command</i> to a JVM instance. See Appendix B for a complete description of mtask commands.
<code>-help</code>	Display usage information and exit.
<code>-host name</code>	The IP address of the host running the server JVM instance. The default is 127.0.0.1.
<code>-killall</code>	Terminate all launched JVM instances.
<code>-killserver</code>	Terminate the master JVM server instance.
<code>-port number</code>	The TCP port bound to the server JVM instance. The default port is 7777.
<code>-testingmode testprefix</code>	Equivalent to the TESTING_MODE mtask command described in TABLE B-1.
<code>-warmup</code> <code>-initClasses classListFile</code> <code>-precompileMethods methodListFile</code>	Specify text files containing the warmup lists for class initialization and method precompilation. See Appendix B for a description of the warmup list syntax.

Java ME System Properties

In addition to the standard Java SE system properties, CDC supports the standard Java ME system properties supported by CLDC 1.1 and MIDP 2.0. These system properties are described in [TABLE D-1](#).

TABLE D-1 CDC System Properties

System Property	Default Value	Description
microedition.commports	No default	Comma-delimited list of available communications ports
microedition.configuration	cdc	Java ME configuration
microedition.encoding	ISO_LATIN_1	Unicode character encoding
microedition.hostname	No default	Host platform
microedition.locale	en-US	System locale
microedition.platform	j2me	Java platform
microeditionprofiles	No default	Java ME profile
microedition.securerandom.nofallback	false	Disable the mechanism that allows the CDC Java runtime environment to fallback to using /dev/urandom if /dev/random doesn't have enough entropy to work properly. See Section 6.2.4, "Seed Generation for Random Number Generation" on page 6-5 for more information.
cdcams.decorations	false	Display native window decorations.
cdcams.presentation	No default	Top-level presentation mode class.

TABLE D-1 CDC System Properties (*Continued*)

System Property	Default Value	Description
cdcams.repository	CVMHOME/ repository	Location of application repository.
cdcams.verbose	false	Display extra diagnostic information.
java.ext.dirs	CVMHOME/ lib	Specifies one or more directories to search for installed optional packages, each separated by File.pathSeparatorChar.

For a list of the standard Java SE system properties, see the description of `java.lang.System.getProperties()` in the CDC specification.

Serial Port Configuration Notes

The `javax.microedition.io.CommConnection` interface allows a CDC Java runtime environment to expose an OS-level serial port as a logical serial port connection. This appendix shows how to configure an OS-level serial port on a Linux system so that a Java application can access the corresponding logical serial port connection.

Note – While this example is based on the RS-232 serial interface implementation of `CommConnection` in `com.sun.cdc.io.j2me.comm.Protocol`, an alternate implementation could use the `CommConnection` interface to support other forms of serial communication such as IrDA.

TABLE E-1 Serial Communications References

Interface	Document
RS-232 serial communications	http://www.tldp.org/HOWTO/Serial-HOWTO-4.html
minicom serial communications program	<code>minicom(1)</code>
Serial port configuration	<code>setserialport(8)</code>
Serial port driver interface	<code>ttyS(4)</code>

E.1 Serial Port Setup

1. Setup a serial cable connection between two Linux computers.

Become super-user.

```
% su  
#
```

This step is necessary to allow non-root users to access the serial port.

2. Configure the serial port to use IRQ 4.

```
# setserial /dev/ttys0 irq 4
```

3. Change the file access permissions for the serial port and the lock file.

```
# chmod 777 /dev/ttys0 /var/lock
```

This allows other users to access the serial port.

4. Launch the minicom(1) serial communications program in setup mode.

```
# minicom -s
```

a. Select Serial port setup from the [configuration] menu.

b. In the setup menu, type A to change the Serial Device setting.

If the Serial Device setting is /dev/modem, then change it to /dev/ttys0.

c. Press <ENTER> to confirm the change.

d. Press <ENTER> again to exit the setup menu.

e. Select the Save setup as dfl menu option.

f. Select the Exit menu option.

This will initialize the serial port.

g. Type <CONTROL>-a q to finally exit minicom(1)

5. Follow a similar configuration procedure with the other computer connected to the serial cable.

E.2 OS-Level Testing

The serial connection between the two computers can be tested with the `minicom(1)` serial communications program.

- 1. Remotely login to each computer.**
- 2. Launch the `minicom(1)` serial communications program on each computer.**
- 3. Type some text into one of the `minicom(1)` windows.**
- 4. Type `<CONTROL>-a q` to finally exit `minicom(1)`.**

This should determine that the serial connection is correct.

Platform Font Administration Notes

Note – These notes show how to install and configure some TrueType fonts for use by the Xft font server running on the Suse Linux 9.1 platform. Next, these fonts are mapped to the Qt logical font names that are used by the CDC Java runtime environment. Other target platforms will vary greatly.

Suse Linux uses the XFree86 server to provide an X11-based desktop infrastructure for supporting graphical desktop environments like KDE and Gnome. X11 font administration has evolved to support remote font servers for international fonts as well as font matching and substitution that allows a single logical font to use multiple physical fonts for supporting different character set ranges.

TABLE F-1 Font Management References

Title	URL	Description
<i>Qt font support</i>	http://trolltech.com	The Qt source release includes the <code>qtnconfig</code> configuration utility which has online help that describes its options and features.
<i>font-config</i>	http://fontconfig.org	A library for font customization and configuration.
<i>Fonts in XFree86</i>	http://www.xfree86.org/4.4.0/fonts.html	Font support in XFree86 and the Xft font system.

This example of installing and configuring a new TrueType font is based installing and configuring the Thorndale font family contained in the files `thornbcp.ttf`, `thorncp.ttf`, `thornicp.ttf` and `thornzcp.ttf`. These font files must be acquired separately.

1. **Install the TrueType fonts in a personal font directory in `$HOME/.fonts` for use by the Xft font server.**

a. If necessary create the personal font directory.

```
% mkdir $HOME/.fonts
```

b. Copy the TrueType font files into the personal font directory.

```
% cp *.ttf $HOME/.fonts
```

c. Change the shell's current directory to the personal font directory.

```
% cd $HOME/.fonts
```

d. Create an index of scalable font files.

```
% mkfontscale
```

e. Create an index of X11 font files.

```
% mkfontdir
```

f. Restart the X11 server.

Logout from the desktop session and login again.

g. Test the X11 font installation with xlsfonts(1) and xfd(1).

```
% xlsfonts | grep -i thorn
...
% xfd -fn "-monotype-thorndale-bold-r-normal--0-0-0-0-p-0-iso8859-1"
```

2. Setup the fonts.conf font configuration system to manage the new fonts.

a. Create a personal fonts.conf configuration file in \$HOME/.fonts.conf.

```
<?xml version="1.0"?>
<!DOCTYPE fontconfig SYSTEM "fonts.dtd">
<fontconfig>
<dir>~/.fonts</dir>
<alias>
  <family>sans-serif</family>
  <prefer><family>Andale Mono</family></prefer>
</alias>
<alias>
  <family>serif</family>
  <prefer><family>Thorndale</family></prefer>
</alias>
<alias>
  <family>Courier</family>
  <prefer><family>Andale Mono</family></prefer>
</alias>
<alias>
  <family>Monospace</family>
  <prefer><family>Andale Mono</family></prefer>
</alias>
</fontconfig>
```

b. Update the fonts.conf cache files.

```
% fc-cache -f
```

c. Determine that the fonts have been configured to match to the fonts.conf logical font names.

```
% fc-list | grep -i thorn
% fc-match serif
```

d. Visually test the font configuration with a non-Qt application like gnome-terminal(1).

3. Check to see if Qt has any extra logical font substitutions that can interfere with the intended font mapping in \$HOME/.fonts.conf.

a. Launch qtconfig.

b. Select the Fonts panel.

c. Remove any extra logical font substitutions for Sans, Sans Serif, Serif, Courier and Monospace.

Extra font substitutions can interfere with the Qt logical font matching mechanism.

d. Visually test the font configuration with a Qt application like the hello example program in the Qt source bundle.

4. Test whether the font mapping between logical Java fonts and logical platform (Qt) fonts is accurate with the following application:

```
import java.awt.*;

class MyFontTest extends Frame {
    MyFontTest() {
        super("MyFontTest");
        setSize(200, 200);
        show();
    }
    public void paint(Graphics g) {
        Font f = new Font("Serif", Font.PLAIN, 12);
        g.setFont(f);
        g.drawString("test string", 50, 50);
    }
    static public void main(String[] args) {
        new MyFontTest();
    }
}
```


Zaurus Installation Notes

This appendix describes how to install the CDC Java runtime environment on a Zaurus personal mobile tool.

G.1 Zaurus System Requirements

The Zaurus personal mobile tool is a PDA based on the Linux operating system, the ARM processor and the Qtopia application environment. It is available mainly for the Japanese market but Dynamism (www.dynamism.com) localizes and markets these devices for the English-speaking market.

TABLE G-1 describes the basic system requirements for a Zaurus personal mobile tool to run the CDC Java runtime environment and CDC AMS.

TABLE G-1 Zaurus System Requirements

Category	Requirement
ROM level	1.41 or greater
Memory	64MB system memory and 256MB SD card memory
Network connection	wireless CF card or Ethernet CF card
Locale	en (English)

TABLE G-2 describes some useful resources for working with the Zaurus personal mobile tool in a development environment.

TABLE G-2 Useful Zaurus Resources

Resource	Purpose
ssh server	Remote login server daemon. An OpenSSH-based server for the Zaurus personal mobile tool is available from www.handango.com .
unzip utility	Unzipping installation bundles. See www.elsix.org .
wireless CF card or Ethernet CF card	OTA provisioning.
SD card	Local software installation and swap space.
SD card reader	Local software installation.

G.2 Installation Procedure

The installation procedure below assumes that the Zaurus already has `ssh` and `unzip` installed on the device.

- 1. Establish a network connection for the Zaurus personal mobile tool.**
- 2. Copy the runtime installation bundle onto an SD card using an SD card reader.**
See **TABLE 2-2** for a description of the installation bundles. For the purposes of these instructions, the installation bundle will be named `runtime.zip`.
- 3. Copy the Qtopia application resources in the resources directory of the documentation bundle onto the SD card.**

`launchams`
`makeswap`
`appmanager.desktop`
`ams.directory`
`28x283DJavaIconB.png`

- 4. Insert the SD card into the Zaurus personal mobile tool.**
- 5. Remote login to the Zaurus personal mobile tool with an ssh client using the root account.**
`% ssh zaurus -l root`
- 6. Change the shell's current directory to the SD card.**
`# cd /mnt/card`

7. Create a Qtopia application directory for CDC AMS.

```
# mkdir /home/QtPalmtop/apps/Java
```

8. Copy the Qtopia application resource files from the card to the Qtopia application directory.

```
# cp ams.directory /home/QtPalmtop/apps/Java/.directory
# chmod 777 /home/QtPalmtop/apps/Java/.directory
# cp appmanager.desktop /home/QtPalmtop/apps/Java
# chmod 777 /home/QtPalmtop/apps/Java/appmanager.desktop
```

9. Copy the CDC AMS icon to the Qtopia icon directories.

```
# cp 28x283DJavaIconB.png /home/QtPalmtop/pics
# cp 28x283DJavaIconB.png /home/QtPalmtop/pics144
```

10. Install the CDC AMS driver script.

```
# cp launchams /home/QtPalmtop/bin
# chmod 755 /home/QtPalmtop/launchams
```

[CODE EXAMPLE G-1](#) contains the text for the launchams driver script.

11. Create a swap file on the SD card.

```
# cd /mnt/card
# sh < /home/zaurus/swap.sh
```

12. Make a directory to hold the CDC Java runtime environment.

```
# mkdir /home/cdcams
```

Because the Personal Profile version of CDC AMS includes a simple user account system and other system features that require writing to local files, the CDC Java runtime environment should be installed on a read-write partition.

13. Edit the `launchams` script and change the definition of the `CDCAMS` variable to refer to the directory containing the CDC Java runtime environment.

```
# vi /home/QtPalmtop/bin/launchams
```

14. Copy the runtime installation bundle from the SD card onto the device.

```
# cp runtime.zip /home/cdcams
```

15. Unpack the installation bundle.

```
# cd /home/cdcams
# unzip runtime.zip
```

The CDC Java runtime environment should now be installed on the Zaurus device. It can be run using the same procedure described in [Chapter 5](#).

CODE EXAMPLE G-1 launchams Script for Zaurus

```
#!/bin/sh
#
# use cvmc to launch the Personal Profile version
# of the CDC AMS appmanager in a client JVM instance
#
CDCAMS=/home/cdcams
SERVERLOG=/tmp/ams_server.log
MANAGERLOG=/tmp/ams_manager.log

AMSCLIENTJAR=$CDCAMS/lib/appmanager-client.jar
AMSMANAGERJAR=$CDCAMS/lib/appmanager.jar
AMSCLASS=com.sun.appmanager.impl.CDCAmsAppManager
PMODEJAR=$CDCAMS/lib/AwtPDA_PresentationMode.jar
PMODECLASS=com.sun.appmanager.impl.presentation.AwtPDA.AwtPDAPresentationMode
XMLJAR=$CDCAMS/lib/j2me_xml_cdc.jar
CLASSLIST=$CDCAMS/repository/profiles/classesList.txt
METHODLIST=$CDCAMS/repository/profiles/methodsList.txt

#
# launch the mtask server
#
$CDCAMS/bin/cvm \ 
-Xbootclasspath/a:$AMSCLIENTJAR \
-Xserver:port=7788,initClasses=$CLASSLIST,precompileMethods=$METHODLIST \
> $SERVERLOG 2>&1 &

#
# wait for the mtask server to finish launching before cloning it
#
sleep 5

#
# use the cvmc driver utility
# to launch the CDC AMS implementation
# in a client JVM instance
#
$CDCAMS/bin/cvmc \
-host 127.0.0.1 \
-port 7788 \
-command JDETACH \
-XappName=$0 \
-Xbootclasspath/a:$AMSMANAGERJAR:$PMODEJAR:$XMLJAR \
-Dcdcams.presentation=$PMODECLASS \
$AMSCLASS \
-port 7788 \
-server 127.0.0.1 \
> $MANAGERLOG 2>&1
```

Cobalt Installation Notes

This appendix describes how to install the CDC Java runtime environment on a Cobalt Qube.

The Cobalt Qube 2 is a Linux/MIPS-based server. See <http://www.linux-mips.org/wiki/index.php/Cobalt> for more information about the Qube 2.

H.1 Installation Procedure

1. **Remote login to the Cobalt Qube with an ssh client using the root account.**

```
% ssh qube -l root
```

2. **Copy the CDC AMS driver script from the resources directory of the documentation bundle.**

```
# cp launchams /usr/bin  
# chmod 755 /usr/bin/launchams
```

3. **Make a directory to hold the CDC Java runtime environment.**

```
# mkdir /home/cdcams
```

4. **Copy the runtime installation bundle into the new directory.**

```
# cp runtime.zip /home/cdcams
```

See [TABLE 2-2](#) for a description of the installation bundles. For the purposes of these instructions, the installation bundle will be named *runtime.zip*.

5. **Unpack the installation bundle.**

```
# cd /home/cdcams  
# unzip runtime.zip
```

The CDC Java runtime environment should now be installed on the Cobalt Qube. It can be run using the same procedure described in [Chapter 5](#).

CODE EXAMPLE H-1 launchams Script for the Cobalt Qube

```
#!/bin/sh
#
# use cvmc to launch the Personal Basis Profile version
# of the CDC AMS appmanager in a client JVM instance
#
CDCAMS=/home/cdcams
SERVERLOG=/tmp/ams_server.log
MANAGERLOG=/tmp/ams_manager.log

AMSCLIENTJAR=$CDCAMS/lib/appmanager-client.jar
AMSMANAGERJAR=$CDCAMS/lib/appmanager.jar
PMODEJAR=$CDCAMS/lib/PBP_PresentationMode.jar
PMODECLASS=com.sun.appmanager.impl.presentation.PBP.PBPPresentationMode
AMSCLASS=com.sun.appmanager.impl.CDCAmsAppManager
CLASSLIST=$CDCAMS/repository/profiles/classesList.txt
METHODLIST=$CDCAMS/repository/profiles/methodsList.txt
PBP_SCREEN_BOUNDS=0,342-640x100
export PBP_SCREEN_BOUNDS

#
# launch the mtask server
#
$CDCAMS/bin/cvm \
-Xbootclasspath/a:$AMSCLIENTJAR \
-Xserver:port=7788,initClasses=$CLASSLIST,precompileMethods=$METHODLIST \
> $SERVERLOG 2>&1 &

#
# wait for the mtask server to finish launching before cloning it
#
sleep 5

#
# use the cvmc driver utility
# to launch the CDC AMS implementation
# in a client JVM instance
#
$CDCAMS/bin/cvmc \
-host 127.0.0.1 \
-port 7788 \
-command JDETACH \
-Xbootclasspath/a:$AMSMANAGERJAR:$PMODEJAR \
-Dcdcams.presentation=$PMODECLASS \
$AMSCLASS \
-port 7788 \
-server 127.0.0.1
> $MANAGERLOG 2>&1
```

Provisioning Server Notes

Note – The OTA provisioning feature of CDC AMS requires integration of the J2EE Client Provisioning RI (CPRI) source code into the CDC build system. This feature cannot be enabled directly with the binary implementation of CDC AMS. Instead, the CDC build system must be first configured with the CPRI source bundle and then used to build a client version of CDC AMS.

Note – The Personal Profile version of CDC AMS includes the provisioning feature to discover and download various types of content, including dynamic content like xlets. This feature is based on the `com.sun.appmanager.ota` package which can support a variety of provisioning technologies. The Personal Profile version of CDC AMS includes a specific example implementation based on the the J2EE Client Provisioning Server Reference Implementation (CPRI). The Personal Basis Profile version of CDC AMS does not include an equivalent provisioning feature.

This appendix describes how to setup a CPRI server for use with the Personal Profile version of CDC AMS. This includes the following steps:

- [Download the J2EE 1.3.1 SDK](#)
- [J2EE Server Setup](#)
- [Download the CPRI Server](#)
- [CPRI Server Setup](#)
- [Testing](#)

I.1

Download the J2EE 1.3.1 SDK

The J2EE 1.3.1 SDK is available from:

<http://java.sun.com/j2ee/1.3/download.html>

I.2

J2EE Server Setup

The J2EE 1.3.1 SDK download page includes standard installation and setup instructions for different target platforms. The following notes provide a brief summary of the J2EE server installation process.

1. Add the following permissions to the server.policy file before starting the J2EE server.

Using a text editor, edit the file \${J2EE_HOME}/lib/security/server.policy and append the following lines:

CODE EXAMPLE I-1 Policy Statement for server.policy File

```
grant {  
    permission java.io.FilePermission  
    "${user.home}${/}MyRepository${/}-",  
    "read,write,delete";  
    permission java.io.FilePermission  
    "${user.home}${/}MyRepository",  
    "read,write,delete";  
};
```

These grant entries enable servlets to read, write and delete local files and allow the `ri-test` servlet to create a download repository in the home directory of the user running the J2EE server.

2. Set the JAVA_HOME, J2EE_HOME and PATH environment variables.

```
# set JAVA_HOME=<path to j2sdk1.4.2>  
# set J2EE_HOME=<path to j2sdkee1.3.1>  
# set PATH=${PATH}: ${J2EE_HOME}/bin  
# export JAVA_HOME J2EE_HOME PATH
```

3. Start the Cloudscape database.

```
# cloudscape -start &
```

Wait for the Cloudscape database to initialize before starting the J2EE server.

4. Start the J2EE server.

```
# j2ee -verbose &
```

Diagnostic messages will be displayed on `stdout` from the cloudscape database and the server as they start. The server will listen on port 1050, provide web service on port 8000 and secure web service at port 7000.

I.3

Download the CPRI Server

The CPRI server is available at:

<http://java.sun.com/j2ee/provisioning/download.html>

I.4

CPRI Server Setup

The CPRI server download page includes standard installation and setup instructions. The following notes provide a brief summary of the CPRI installation process.

In this procedure, we repackage the CPRI classes in their own EAR file to include OMA-aware adapter classes that can communicate with the appmanager client. This is in turn deployed to the J2EE server to create the *ri-test* servlet.

1. **Set the J2EE_HOME and JSR124_HOME variables to specify the locations of the J2EE and CPRI servers.**

```
# J2EE_HOME=location/j2sdkee1.3.1 ; export J2EE_HOME  
# JSR124_HOME=location/j2ee_cp_ri_1_0 ; export JSR124_HOME
```

2. **Build an EAR file using the procedure described in the *CDC Build System Guide*.**
3. **Deploy the EAR file for the CPRI server.**

```
# java -classpath ${J2EE_HOME}/lib/j2ee.jar:${J2EE_HOME}/lib/locale\  
-Dorg.omg.CORBA.ORBInitialPort=1050 \  
-Dcom.sun.enterprise.home=${J2EE_HOME} \  
-Djava.security.policy=${J2EE_HOME}/lib/security/server.policy \  
com.sun.enterprise.tools.deployment.main.Main \  
-deploy earfile local
```

When this step is completed, the CPRI server should be deployed and running on the J2EE server. It will be accessible through a web browser at the following URL:

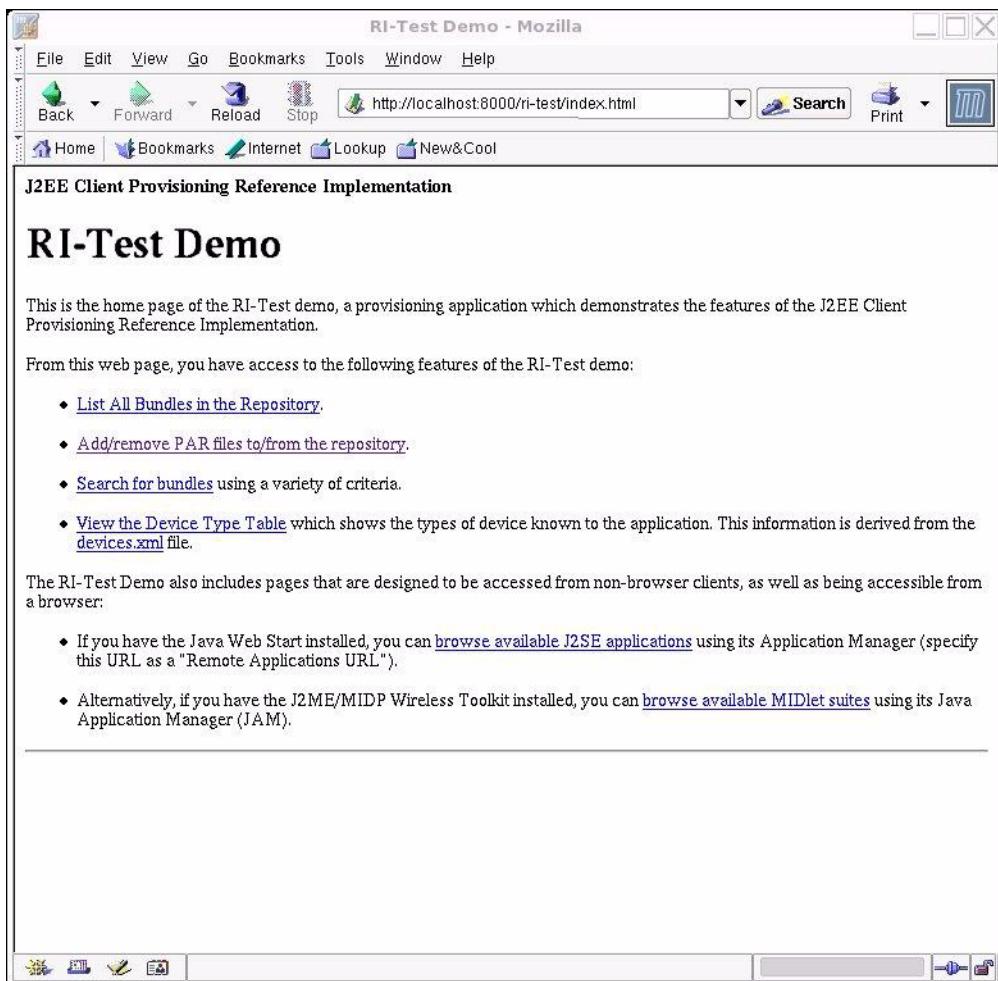
<http://j2eeserver:8000/ri-test/index.html>

Once the *ri-test* servlet is deployed to the J2EE server, you can view *ri-test* page:

<http://j2ee-server:8000/ri-test/index.html>

The `ri-test` page should look something like **FIGURE I-1**:

FIGURE I-1 `ri-test` Page



I.5 Deploy a PAR File

A PAR file is a mechanism for staging downloadable content in a CPRI server.

1. Create a PAR file containing an xlet.

The resources directory in the documentation bundle includes a sample PAR file for a helloworld xlet.

2. Upload the PAR file to the J2EE server.

Select the Add/remove PAR files to/from the repository link. The browser will display the upload page shown in [FIGURE I-2](#).

FIGURE I-2 ri-test Upload Page

Bundle Repository Administration - Mozilla

File Edit View Go Bookmarks Tools Window Help

Back Forward Reload Stop http://localhost:8000/ri-test/admin/stocking Search Print

Home Bookmarks Internet Lookup New&Cool

Bundle Repository Administration

Repository Contents

PAR	Uploaded	Bundles	
5	Thu Nov 04 21:30:50 PST 2004	16	remove

Upload a PAR File

Upload a PAR file by entering its path and clicking "Upload PAR file". For example, click on Browse and go to .../test/apps/test.par.

PAR file: [Browse...](#)

Empty the Repository

If you want to start over, you can empty out the repository by clicking on this button:

[Remove all PAR files from repository](#)

[Back to top](#)

Select the Browse button to navigate to the target PAR file, and then select the Upload PAR file button. The PAR file should now be stored on the CPRI server and available for OTA deployment.

Verify that the PAR file was correctly uploaded by viewing the list by clicking on "List all Bundles in the Repository". Also verify all the contents are displayed correctly at <http://j2ee-server:8000/ri-test/oma>.

I.6 Testing

Test the ri-test server with a CDC AMS client.

When the upload process is complete, the ri-test server is ready for testing.

1. **Launch the Personal Profile version of CDC AMS.**
2. **Set the Discovery URL in System>Preferences>General for locating content.**

`http://J2EE-server:8000/ri-test/oma`

3. **Start the provisioning application at System>App Store.**

The provisioning application will display a set of staged applications for deployment. Selecting an application will cause it to be downloaded and deployed in the application repository.