

**IBM SecureWay Cryptographic Products  
IBM 4758 PCI Cryptographic Coprocessor  
Custom Software Developer's Toolkit Guide**

July 30, 1998

Security Solutions and Technology Department

IBM Corporation  
8501 IBM Drive  
Charlotte, North Carolina 28262-8563

**30-JUL-98, 09:17**

**Note!**

Before using this information and the products it supports, be sure to read the general information under Appendix F, "Notices" on page F-1.

**First Edition (June, 1998)**

Changes are made periodically to the information herein; before using this publication in connection with the operation of IBM systems, consult your IBM representative to be sure you have the latest edition and any Technical Newsletter.

IBM does not stock publications at the address given below; requests for IBM publications should be made to your IBM representative or to the IBM branch office that serves your location.

Reader's comments can be communicated by e-mail to George Dolan, [gmdolan@us.ibm.com](mailto:gmdolan@us.ibm.com), or the comments can be addressed to IBM Corporation, Department VM9A, MG81/204, 8501 IBM Drive, Charlotte, NC 28262-8563, U.S.A. IBM may use or distribute whatever information you supply in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 1998. All rights reserved.**

Note to U.S. Government Users — Documentation related to restricted rights — Use, duplication or disclosure is subject to restrictions set forth in GSA ADP Schedule Contract with IBM Corp.

---

# Table of Contents

<b>About This Book</b> . . . . .	vii
Prerequisite Knowledge . . . . .	vii
Organization of This Book . . . . .	vii
Typographic Conventions . . . . .	viii
Syntax Diagrams . . . . .	viii
Related Publications . . . . .	ix
General Interest . . . . .	ix
CCA Support Program Publications . . . . .	ix
Custom Software Publications . . . . .	ix
Cryptography Publications . . . . .	x
Summary of Changes . . . . .	xi
<b>Chapter 1. Introduction</b> . . . . .	1-1
Available Documentation . . . . .	1-1
Prerequisites . . . . .	1-2
Development Overview . . . . .	1-3
Development Environment Components . . . . .	1-4
Release Components . . . . .	1-5
<b>Chapter 2. Installation and Setup</b> . . . . .	2-1
Installing the Toolkit . . . . .	2-1
Directories and Files . . . . .	2-1
Preparing the Development Platform . . . . .	2-5
<b>Chapter 3. Developing and Debugging an SCC Application</b> . . . . .	3-1
Development Process Road Map . . . . .	3-1
Special Coding Requirements During Development . . . . .	3-3
Developer Identifiers . . . . .	3-3
Attaching with the Debugger . . . . .	3-3
Compiling, Assembling, and Linking . . . . .	3-3
CP/Q Base Operating System Function Support . . . . .	3-4
C Run-Time Library Support . . . . .	3-4
Supported Functions and Global Variables . . . . .	3-4
Unsupported Functions and Global Variables . . . . .	3-5
Compiler Options . . . . .	3-6
VisualAge C++ (VACPP) Options . . . . .	3-6
Microsoft Visual C++ (MSVC++) Options . . . . .	3-7
Assembler Options . . . . .	3-8
Linker Options . . . . .	3-9
ILINK (VACPP Linker) . . . . .	3-9
LINK (MSVC++ Linker) . . . . .	3-9
Librarian Options . . . . .	3-10
Translating . . . . .	3-10
Building Read-Only Disk Images . . . . .	3-10
Downloading and Debugging . . . . .	3-11
<b>Chapter 4. Testing an SCC Application in a Production Environment</b> . . . . .	4-1
<b>Chapter 5. Packaging and Releasing an SCC Application</b> . . . . .	5-1

<b>Appendix A. An Overview of the Development Process</b> . . . . .	A-1
<b>Appendix B. Using CLU</b> . . . . .	B-1
<b>Appendix C. How to Reboot the IBM 4758</b> . . . . .	C-1
<b>Appendix D. Building SCC Applications with Microsoft Developer Studio</b>	
<b>97</b> . . . . .	D-1
Required Settings for the Host-Side Portion of an SCC Application . . . . .	D-1
Required Settings for the Coprocessor-Side Portion of an SCC Application . . . . .	D-1
<b>Appendix E. Using Signer and Packager</b> . . . . .	E-1
Coprocessor Memory Segments and Security . . . . .	E-1
The Signer Utility (TKNSGMR.EXE) . . . . .	E-4
Signer Operations . . . . .	E-5
Signer Cryptographic Functions . . . . .	E-5
Signer Miniboot Command Functions . . . . .	E-5
Signer Miscellaneous Functions . . . . .	E-5
Signer IBM-Specific Functions . . . . .	E-5
EMBURN2 - Load Software into Segment 2 . . . . .	E-6
EMBURN3 - Load Software into Segment 3 . . . . .	E-7
ESIG3 - Build Emergency Signature for Segment 3 . . . . .	E-8
ESTOWN3 - Establish Ownership of Segment 3 . . . . .	E-9
HASH_GEN - Generate Hash for File . . . . .	E-10
HASH_VER - Verify Hash of File . . . . .	E-10
KEYGEN - Generate RSA Key Pair . . . . .	E-10
REMBURN2 - Replace Software in Segment 2 . . . . .	E-11
REMBURN3 - Replace Software in Segment 3 . . . . .	E-12
SUROWN2 - Surrender Ownership of Segment 2 . . . . .	E-13
SUROWN3 - Surrender Ownership of Segment 3 . . . . .	E-14
File Description Arguments . . . . .	E-15
Signature Key Arguments . . . . .	E-15
Image File Arguments . . . . .	E-16
Trust and Countersignature Arguments . . . . .	E-16
Targeting Arguments . . . . .	E-17
The Packager Utility (TKNPKGR.EXE) . . . . .	E-21
<b>Appendix F. Notices</b> . . . . .	F-1
Copying and Distributing Softcopy Files . . . . .	F-1
Trademarks . . . . .	F-2
<b>List of Abbreviations and Acronyms</b> . . . . .	X-1
<b>Glossary</b> . . . . .	X-3
<b>Index</b> . . . . .	X-9

---

## Figures

1-1.	Development Process Overview	1-3
2-1.	Toolkit Directory Structure	2-2
3-1.	Development Process Road Map	3-2
E-1.	State Transitions for Segment 2	E-3
E-2.	State Transitions for Segment 3	E-4



---

## About This Book

The *IBM 4758 PCI Cryptographic Coprocessor Custom Software Developer's Toolkit Guide* describes the Developer's Toolkit and its components, including the tools that enable developers to:

- Build applications for the IBM 4758 PCI Cryptographic Coprocessor
- Load applications under development into a coprocessor
- Debug applications under development running within a coprocessor

The primary audience for this book are developers who are creating applications to use with the coprocessor. People who are interested in packaging, distribution, and security issues for custom software should also read this book.

---

## Prerequisite Knowledge

The reader of this book should understand how to perform basic tasks (including editing, system configuration, file system navigation, and creating application programs) on the host machine. Familiarity with the coprocessor hardware (as described in the *IBM 4758 PCI Cryptographic Coprocessor Technical Overview for Original Equipment Manufacturers*), the CP/Q++ operating system that runs within the coprocessor (as described in the *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System Overview*), and the use of the IBM's Common Cryptographic Architecture (CCA) application and support program (as described in the *IBM 4758 PCI Cryptographic Coprocessor CCA Support Program, SC31-8610*) may also be helpful.

People who are interested in packaging, distribution, and security issues for custom software will need to understand the use of the CCA Support Program and should be familiar with the coprocessor's security architecture as described in IBM Research Report RC21102, "Building a High-Performance, Programmable Secure Coprocessor." See "Cryptography Publications" on page x for information on how to obtain this research report.

---

## Organization of This Book

This book is organized as follows:

Chapter 1, "Introduction" describes the documentation available to a developer of an SCC application, lists the prerequisites for development, describes the development process, and lists the tools used during development.

Chapter 2, "Installation and Setup" describes how to install the Developer's Toolkit and how to prepare an IBM 4758 PCI cryptographic coprocessor for use as a development platform.

Chapter 3, "Developing and Debugging an SCC Application" discusses in detail the use of each of the tools used during development of an SCC application.

Chapter 4, "Testing an SCC Application in a Production Environment" describes how to load production-level software into the coprocessor used as a development platform.

Chapter 5, “Packaging and Releasing an SCC Application” describes how to prepare an SCC application to be distributed to end users.

Appendix A, “An Overview of the Development Process” lists the steps a developer needs to perform during development and testing of an SCC application.

Appendix B, “Using CLU” briefly describes the use of the Coprocessor Load Utility.

Appendix C, “How to Reboot the IBM 4758” describes several ways to reboot a cryptographic coprocessor. If an application has been loaded into the coprocessor, it starts to run after the reboot is complete.

Appendix D, “Building SCC Applications with Microsoft Developer Studio 97” describes how to configure Microsoft Developer Studio 97\*\* to ensure the proper compiler and linker options are used to build an SCC application.

Appendix E, “Using Signer and Packager” describes the use of the signer and packager utilities and explains why the design of the coprocessor makes these utilities necessary.

Appendix F, “Notices” includes product and publication notices.

A list of abbreviations, a glossary, and an index complete the manual.

---

## Typographic Conventions

This publication uses the following typographic conventions:

- Commands that you enter verbatim onto the command line are presented in **bold** type.
- Variable information and parameters, such as file names, are presented in *italic* type.
- The names of items that are displayed in graphical user interface (GUI) applications—such as pull-down menus, checkboxes, radio buttons, and fields—are presented in **bold** type.
- Items displayed within pull-down menus are presented in ***bold italic*** type.
- System responses in a non-GUI environment are presented in monospace type.
- Web addresses and directory paths are presented in *italic* type.

---

## Syntax Diagrams

The syntax diagrams in this section follow the typographic conventions listed in “Typographic Conventions” described previously. Optional items appear in brackets. Lists from which a selection must be made appear in braces with vertical bars separating the choices. See the following example.

**COMMAND** *firstarg* [*secondarg*] {**a** | **b**}

A value for *firstarg* must be specified. *secondarg* may be omitted. Either **a** or **b** must be specified.



---

## Related Publications

Many of the publications listed below under “General Interest,” “CCA Support Program Publications,” and “Custom Software Publications” are available in Adobe Acrobat\*\* portable document format (PDF) at <http://www.ibm.com/security/cryptocards>.

### General Interest

The following publications may be of interest to anyone who needs to install, use, or write applications for a PCI Cryptographic Coprocessor:

- *IBM 4758 PCI Cryptographic Coprocessor General Information Manual*, GC31-8608 (version -01 or later)
- *IBM 4758 PCI Cryptographic Coprocessor Installation Manual*, SC31-8623
- *IBM 4758 PCI Cryptographic Coprocessor Custom Software Installation Manual*

### CCA Support Program Publications

The following publications may be of interest to readers who intend to use a PCI Cryptographic Coprocessor to run IBM’s Common Cryptographic Architecture (CCA) Support Program:

- *IBM 4758 PCI Cryptographic Coprocessor CCA Support Program*, SC31-8610
- *IBM 4758 CCA Basic Services Reference and Guide*, SC31-8609

### Custom Software Publications

The following publications may be of interest to persons who intend to write applications or operating systems that will run on a PCI Cryptographic Coprocessor:

- *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System Overview*
- *IBM 4758 PCI Cryptographic Coprocessor Custom Software Interface Reference*
- *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System Application Programming Reference*
- *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System C Runtime Library Reference*
- *IBM 4758 PCI Cryptographic Coprocessor CCA User Defined Extensions Programming Reference*
- *IBM 4758 PCI Cryptographic Coprocessor Interactive Code Analysis Tool (ICAT) User’s Guide*
- *AMCC S5933 PCI Controller Data Book*, available from Applied Micro Circuits Corporation, 6290 Sequence Drive, San Diego, CA 92121-4358. Phone 1-800-755-2622 or 1-619-450-9333. The manual is available online as an Adobe Acrobat\*\* PDF file at <http://www.amcc.com/pdfs/5933db.pdf>.

## Cryptography Publications

The following publications describe cryptographic standards, research, and practices applicable to the PCI Cryptographic Coprocessor:

- “Building a High-Performance, Programmable Secure Coprocessor,” S.W. Smith and S.H. Weingart, Research Report RC21102, IBM T.J. Watson Research Center, February 1998.
- “Using a High-Performance, Programmable Secure Coprocessor, S.W. Smith, E.R. Palmer, and S.H. Weingart, in *FC98: Proceedings of the Second International Conference on Financial Cryptography*, Anguilla, February 1998. To appear, Springer-Verlag LNCS, 1998.
- “Secure Coprocessing Research and Application Issues,” S.W. Smith, Los Alamos Unclassified Release LA-UR-96-2805, Los Alamos National Laboratory, August 1996.
- “Secure Coprocessing in Electronic Commerce Applications,” B.S. Yee and J.D. Tygar, in *Proceedings of the First USENIX Workshop on Electronic Commerce*, New York, July 1995.
- “Transaction Security Systems,” D.G. Abraham, G.M. Dolan, G.P. Double, and J.V. Stevens, in *IBM Systems Journal* Vol. 30 No. 2, 1991, G321-0103.
- “Trusting Trusted Hardware: Towards a Formal Model for Programmable Secure Coprocessors,” S.W. Smith and V. Austel, in *Proceedings of the Third USENIX Workshop on Electronic Commerce*, Boston, August 1998.
- “Using Secure Coprocessors,” B.S. Yee (Ph.D. thesis), Computer Science Technical Report CMU-CS-94-149, Carnegie-Mellon University, May 1994.
- *IBM Systems Journal* Vol. 32 No. 3, 1993, G321-5521
- *IBM Journal of Research and Development* Vol. 38 No. 2, 1994, G322-0191
- *Applied Cryptography: Protocols, Algorithms, and Source Code in C, Second Edition*, Bruce Schneier, John Wiley & Sons, Inc. ISBN 0-471-12845-7 or ISBN 0-471-11709-9
- *ANSI X9.31 Public Key Cryptography Using Reversible Algorithms for the Financial Services Industry*
- *Internet Engineering Taskforce RFC 1321*, April 1992, MD5
- *ISO 9796 Digital Signal Standard*
- *Secure Electronic Transaction Protocol Version 1.0*, May 31, 1997
- *USA Federal Information Processing Standard (FIPS):*
  - *Data Encryption Standard*, 46-1-1988
  - *Secure Hash Algorithm*, 180-1, May 31, 1994
  - *Cryptographic Module Security*, 140-1
- *Derived Test Requirements for FIPS PUB 140-1*, W. Havener, R. Medlock, L. Mitchell, and R. Walcott. MITRE Corporation, March 1995.

IBM Research Reports can be obtained from:

IBM T.J. Watson Research Center  
Publications Office, 16-220  
P.O. Box 218  
Yorktown Heights, NY 10598

Back issues of the *IBM Systems Journal* and the *IBM Journal of Research and Development* may be ordered by calling (914) 945-3836.

---

## Summary of Changes

This first edition of the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Developer's Toolkit Guide* contains product information that is current with the IBM 4758 PCI Cryptographic Coprocessor announcements made through May, 1998.



---

## Chapter 1. Introduction

The Developer's Toolkit is a set of libraries, include files, and utility programs that help a developer build, load, and debug applications written in C or assembler for the IBM 4758 PCI Cryptographic Coprocessor. An application that runs within the coprocessor is known as an "agent" or an "SCC application".<sup>1</sup>

The Developer's Toolkit, a commercial compiler and linker, and a PC running Windows NT constitute a complete development environment for the IBM 4758. IBM's CCA Support Program feature is required in order to create a version of an application suitable for distribution. This chapter includes:

- A description of the documentation available to a developer of an SCC application and suggestions on the order in which the introductory material should be read
- A list of hardware and software necessary to develop and release SCC applications
- An overview of the development process
- A description of the software that constitutes the development environment
- A description of the software used to prepare an SCC application for release

---

### Available Documentation

"Related Publications" on page ix lists over twenty publications, many of which are of particular interest to the developer of an SCC application. It may be helpful to read the following manuals in the order listed prior to starting development:

1. *IBM 4758 PCI Cryptographic Coprocessor Technical Overview for Original Equipment Manufacturers*, which contains background information a reader of this book is assumed to understand.
2. *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System Overview*, which describes the features of the coprocessor operating system.
3. *IBM 4758 PCI Cryptographic Coprocessor General Information Manual*, GC31-8608 which provides a basic understanding of IBM's Common Cryptographic Architecture for the IBM 4758.<sup>2</sup>
4. This book, which describes the overall development process and the tools used in the development process.

During development, the following manuals will be of use:

- The *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System Application Programming Reference*, which describes the function calls supplied by the coprocessor operating system.
- The *IBM 4758 PCI Cryptographic Coprocessor Custom Software Interface Reference*, which describes the function calls supplied by the coprocessor device drivers that manage communication, encryption and decryption, random number generation, nonvolatile memory, and other coprocessor services.
- The *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System C Runtime Library Reference*, which describes the function calls supported by the

---

<sup>1</sup> Secure cryptographic coprocessor (SCC) is an alternate name for the IBM 4758 PCI Cryptographic Coprocessor.

<sup>2</sup> This document will be of particular interest to developers writing user-developed extensions for CCA.

full C run-time library supplied by the general version of CP/Q. The customized version of CP/Q that runs in an IBM 4758 does not support the full C run-time library. See "C Run-Time Library Support" on page 3-4 for details.

- Developers writing extensions for IBM's CCA application will also need the *IBM 4758 PCI Cryptographic Coprocessor CCA User Defined Extensions Programming Reference*.
- *IBM 4758 PCI Cryptographic Coprocessor CCA Support Program*, SC31-8610 which describes how to configure IBM's CCA application, which in turn is used by some of the tools in the Developer's Toolkit.

---

## Prerequisites

Prior to the start of development a developer must obtain and install the following:

1. An IBM 4758 model 001 PCI cryptographic coprocessor. Refer to <http://www.ibm.com/security/cryptocards> for ordering information.

The IBM 4758 should be installed in a host following the instructions in the *IBM 4758 PCI Cryptographic Coprocessor Installation Manual*, SC31-8623, which also lists the hardware and software requirements for the host. For application development, the host must be a PC running Windows NT.

2. One of the supported compilers (IBM VisualAge C++<sup>3</sup> or Microsoft Visual C++) and the associated tools, which should be installed following the instructions provided with the compiler. Only the compiler and linker need be installed; other components (visual build environments, and so on) are not required.
3. The IBM 4758 Application Program Development Toolkit (the Developer's Toolkit), available from IBM (order PRPQ 5799-RHB), which should be installed on the same host as the compiler following the instructions in chapter 2 of this manual.

The Developer's Toolkit includes a device driver for the IBM 4758 which should be installed on the host following the instructions in chapter 2 of the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Installation Manual*.

4. Developers writing extensions to IBM's CCA application will also need the IBM 4758 CCA UDX Application Program Development Toolkit Extension (the UDX Toolkit), available from IBM (order PRPQ 5799-RHA), which should be installed on the same host as the Developer's Toolkit following the instructions in chapter 2 of this manual.

The developer must obtain and install the following to prepare an application for release:

1. IBM's CCA Support Program for Windows NT (feature 4376) and a function control vector permitted by the applicable import or export regulations (feature 5200, 5201, or 5202). Information on ordering these items can be obtained from <http://www.ibm.com/security/cryptocards>.

The CCA Support Program should be installed following the instructions in chapter 3 of the *IBM 4758 PCI Cryptographic Coprocessor CCA Support Program*, SC31-8610.

---

<sup>3</sup> IBM offers VisualAge C++ in several packages. Part number 33H4979 is version 3.5 on CD; part number 33H4980 includes documentation. Upgrades are also available.

## Development Overview

As illustrated in Figure 1-1, an SCC application is compiled and linked in the same manner as a host application, using include and library files customized for the coprocessor environment. The executable is then translated to the format understood by CP/Q++ and is downloaded to the coprocessor.

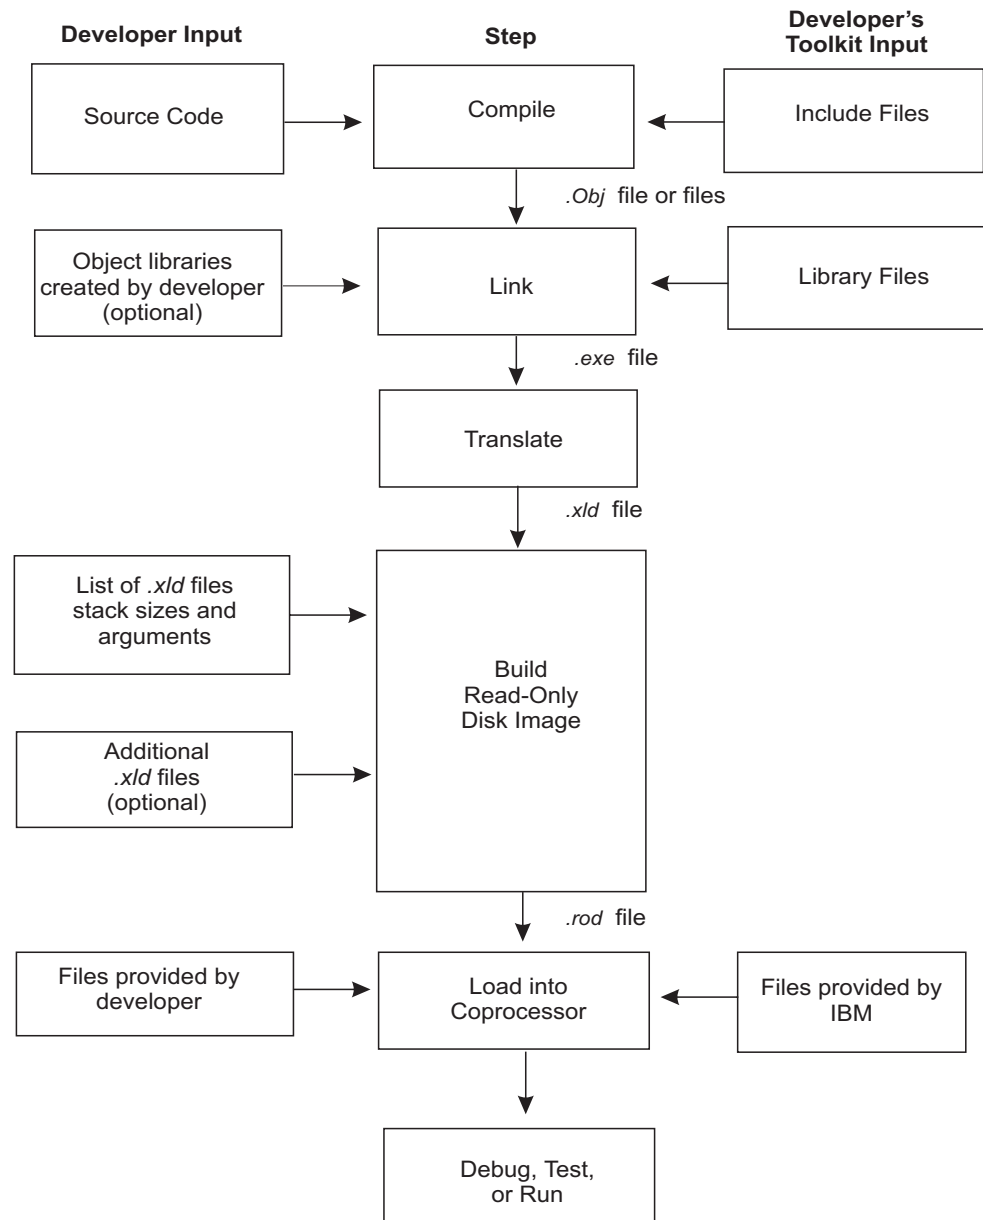


Figure 1-1. Development Process Overview

The following steps are required to build and load SCC applications:

1. Compile the program.
2. Link the program.
3. Translate the executable file into the format required by CP/Q++.
4. Build a read-only disk image.
5. Load the disk image into the coprocessor.

The Developer's Toolkit includes the tools needed to perform steps 3 through 5.

---

## Development Environment Components

The development environment software consists of the following items, most of which are contained in the Developer's Toolkit:

### **Compiler and Linker**

Use either IBM VisualAge\* C++, Microsoft Visual C++\*\*, or—for assembler language code—Microsoft Assembler\*\*. The linker must be compatible with the compiler; for example, use ILINK with VisualAge C++. These are not shipped with the Developer's Toolkit.<sup>4</sup> Part numbers for IBM VisualAge C++ appear in "Prerequisites" on page 1-2.

### **Libraries and Include Files**

Use the Developer's Toolkit libraries (*.lib*) and include files (*.h/.inc*) in place of the libraries and include files shipped with the compiler and assembler. These files furnish the library entry points and routines that SCC applications use to interface with the CP/Q operating system and the CP/Q++ extensions.

### **Utilities**

Use the following utilities to prepare and load SCC applications:

- **Translator:** A program (CPQXLT.EXE) that translates executable (*.exe*) files into the format (*.xld*) required by CP/Q.
- **Disk Builder:** A program (SCCRODSK.EXE) that packages one or more applications into a read-only disk image.
- **Development Reload Utility (DRUID):** A program (DRUID.EXE<sup>5</sup>) that loads an application into a coprocessor configured as a development platform.
- **Coprocessor Load Utility (CLU):** A program (TKNCLU.EXE) that verifies and loads digitally signed system software and coprocessor commands into a coprocessor.

### **CLU Input Files**

The Developer's Toolkit includes several files used as input to CLU during the development process.

### **Debugger**

The IBM Interactive Code Analysis Tool (ICAT) debugger (ICATCPW.EXE) is a Windows NT program that controls and debugs SCC applications.

---

<sup>4</sup> Developers writing extensions for IBM's CCA application must use IBM VisualAge C++.

<sup>5</sup> Development Reload Utility for Insecure Development (DRUID)



### ***Coprocessor Operating System***

The Developer's Toolkit includes two versions of the CP/Q++ embedded operating system:

- A **debug** version (TPRrrss.CLU) used when coding and debugging an application. It contains a debug probe that runs within the coprocessor and services requests from the ICAT debugger.
- A **production** version (TNPrrss.CLU) used to test an application in a production level environment. It does not include the debug probe.

Both versions are supplied as signed disk images that can be loaded into the coprocessor by CLU. They include the CP/Q++ extensions needed to manage the coprocessor hardware, and can include custom extensions specified by the contract between the developer and IBM.

---

## **Release Components**

The software required to prepare an application for release to end users, most of which is contained in the Developer's Toolkit, is listed below.

### ***Utilities***

Use the following utilities to prepare an SCC application for release:

- **Signer:** A program (TKNSGnr.EXE) that generates RSA keypairs and performs other cryptographic operations and that incorporates a read-only disk image into a coprocessor command and digitally signs the command using a developer's private key.
- **Packager:** A program (TKNPKGR.EXE) that combines one or more signed commands into a single file for download to the coprocessor.

### ***CCA Application and Support Program***

Signer and Packager use IBM's Common Cryptographic Architecture (CCA) application to generate digital signatures. The CCA application and support program are not shipped with the Developer's Toolkit. See "Prerequisites" on page 1-2 for more information.



---

## Chapter 2. Installation and Setup

The Developer's Toolkit includes utilities used to build an SCC application and prepare it to be loaded into an IBM 4758 PCI cryptographic coprocessor. This chapter describes how to install the Developer's Toolkit and the UDX Toolkit (if used), discusses the toolkit's directory structure and lists many of the files used during development, and explains how to prepare the coprocessor for use as a development platform.

---

### Installing the Toolkit

The Developer's Toolkit is shipped or made available for download as ZIP files. To install the Developer's Toolkit, unzip each file. Specify the appropriate options to preserve the directory structure in each file and specify the same directory as the target of each unzip command, for example:

```
pkunzip -d scctk1.zip c:  
pkunzip -d scctk2.zip c:
```

The UDX Toolkit is also shipped or made available for download as one or more ZIP files. The UDX Toolkit should be installed in the same manner and into the same target directory as the Developer's Toolkit, for example:

```
pkunzip -d udxtk1.zip c:  
pkunzip -d ukxtk2.zip c:
```

### Directories and Files

The Developer's Toolkit is contained in the directory structure depicted in Figure 2-1 on page 2-2.

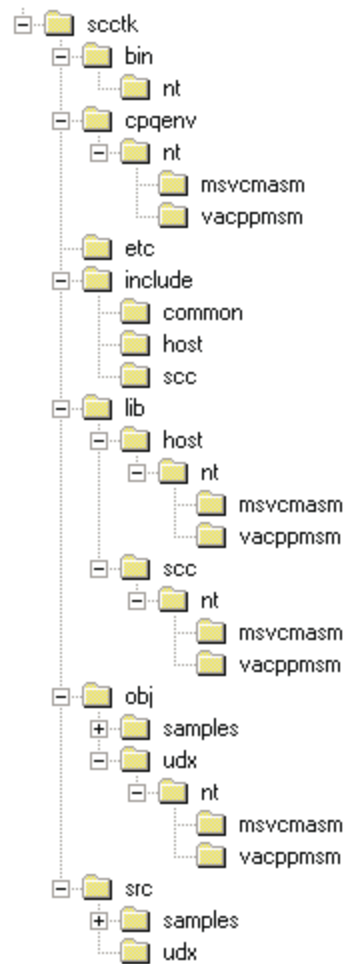


Figure 2-1. Toolkit Directory Structure

- The `scctk\nt\bin` directory contains the following:
  - The NT device driver for the IBM 4758 (CRYPTONT.SYS)
  - The DLL applications on the host use to communicate with applications on the coprocessor (CRYPTONT.DLL)
  - A sample host application that interacts with a “reverse-then-echo” application on the coprocessor (HRE.EXE)
  - The translator utility (CPQXLT.EXE)
  - The read-only disk image builder (SCCRODSK.EXE)
  - The device reload utility (DRUID.EXE)
  - The coprocessor load utility (TKNCLU.EXE)
  - The debugger (ICATCPW.EXE)
  - The signer utility (TKNSGMR.EXE)
  - The packager utility (TKNPKGR.EXE)
  - A utility to display the contents of a read-only disk image file (FMTRODSK.EXE)
  - A utility to display the contents of a CLU file (FMTTKCLU.EXE)
  - DLLs and command files used by the tools in the directory
- The `scctk\cpqenv\nt` directories contain include files (.h and .inc) that define macro variables to customize the include files in the `scctk\include\scc` directory for use with Microsoft Visual C++ (MSVC++) and Microsoft Assembler (MASM)

(the *scctklcpqenv\ntlmsvcmasm* directory) or with IBM VisualAge C++ (VACPP) and MASM (the *scctklcpqenv\ntlvacppmsm* directory).

- The *scctkletc* directory contains files to be used as input to CLU (as described in “Preparing the Development Platform” on page 2-5 and “Downloading and Debugging” on page 3-11), including those listed below. Many files in this directory have names of the form Txxrrrss, where *rrr* indicates the release of the CCA application the file contains or with which the file is associated, *ss* indicates the revision level of the CCA application, and *xx* distinguishes the file from all others.
  - CFCrrrss.CLU, which loads release *rrr* revision *ss* of IBM’s system software into a coprocessor.<sup>1</sup>

CFCrrrss.CLU can only be loaded into an IBM 4758 in the factory-shipped state.<sup>2</sup>
  - TDVrrrss.CLU, which prepares a coprocessor for use as a development platform. TDVrrrss.CLU also loads a “reverse-then-echo” application into the coprocessor so that the developer can verify the development platform is properly configured.<sup>3</sup>

TDVrrrss.CLU can only be loaded into an IBM 4758 that contains release *rrr* revision *ss* of IBM’s system software.<sup>4</sup>

Export regulations may dictate that the version of TDVrrrss.CLU shipped to a particular developer be customized so that the file can only be loaded into a specific coprocessor or a specific set of coprocessors.
  - TPRrrrss.CLU, which loads into a coprocessor a copy of the operating system (CP/Q++) that allows a coprocessor application to be debugged.
 

TPRrrrss.CLU can only be loaded into an IBM 4758 that has been prepared for use as a development platform using TDVrrrss.CLU.<sup>5</sup>
  - TNPrrrss.CLU, which replaces the “debug-enabled” version of CP/Q++ with a production-level copy of CP/Q++ release *rrr* revision *ss*. This allows a developer to test a coprocessor application in a production-level environment.
 

TNPrrrss.CLU can only be loaded into an IBM 4758 that has been prepared for use as a development platform using TDVrrrss.CLU.<sup>5</sup>
  - TL3rrrss.CLU, which clears any state an application under development has saved in nonvolatile memory (so that the application will start next time with a clean slate). TL3rrrss.CLU also loads the “reverse-then-echo” application into the coprocessor.
 

TL3rrrss.CLU can only be loaded into an IBM 4758 that has been prepared for use as a development platform using TDVrrrss.CLU.<sup>5</sup>

---

<sup>1</sup> CFCrrrss.CLU also sets the public key associated with segment 1.

<sup>2</sup> In particular, the public key associated with segment 1 must be the key installed during manufacture.

<sup>3</sup> TDVrrrss.CLU sets the public keys and owner identifiers associated with segments 2 and 3 and loads the reverse-then-echo application into segment 3. Currently, the owner identifier assigned to segment 2 is 3 and the owner identifier assigned to segment 3 is 6.

<sup>4</sup> In particular, segment 2 must be empty and the public key associated with segment 1 must be the key loaded by CFCrrrss.CLU. Loading CCA also causes the key associated with segment 1 to be set to the proper value.

<sup>5</sup> In particular, the public key and owner identifier associated with segment 2 and the image names associated with segments 1 and 2 must have the values CFCrrrss.CLU and TDVrrrss.CLU assign them.

- TR3rrrss.CLU, which reloads the “reverse-then-echo” application into the coprocessor.

TR3rrrss.CLU can only be loaded into an IBM 4758 that has been prepared for use as a development platform using TDVrrrss.CLU.<sup>6</sup>

- S3KCLRPP.DRK and S3KCLRPU.DRK, which contain an RSA keypair and the public key of the keypair, respectively. These files are provided as input to DRUID.
- One or more RSA key token files (file extension .TKN). The developer uses these files with TKNSGNR to generate RSA keys prior to releasing an application.
- TRSrrrss.CLU, which prepares an IBM 4758 that has been used for development to be used in a production setting. TRSrrrss.CLU essentially restores the coprocessor to the state it is in immediately after CFCrrrss.CLU has been loaded.

TRSrrrss.CLU can only be loaded into an IBM 4758 that has been prepared for use as a development platform using TDVrrrss.CLU.<sup>7</sup>

- C2Frrrss.CLU, which restores a coprocessor that has been used for development to the factory-shipped state.

#### Warning

C2Frrrss.CLU updates the public key associated with segment 1. This key can only be updated a few times before the coprocessor runs out of memory in which to store the certificate chain connecting the segment 1 public key to the original key installed at the factory. Users should restore a coprocessor to its factory-shipped state only if absolutely necessary. Note that once the public key associated with segment 1 has been set using CFCrrrss.CLU or by loading IBM's CCA application, it should not be necessary to restore the coprocessor to its factory-shipped state.

- The *scctk\include* directories contain include files (*.h* and *.inc*) that replace the standard include files that ship with MSVCC++, VACPP, and MASM.
  - *scctk\include\common* contains include files that are used to build both SCC applications and host applications that interact with SCC applications.
  - *scctk\include\host* contains include files that are used only to build host applications.
  - *scctk\include\scc* contains include files that are used only to build SCC applications.
- The *scctk\lib* directories contain library (*.lib*) files that augment or replace the standard library files that ship with MSVC++ or VACPP. The *scctk\lib\host* directories contain library files that are used to build host applications and the *scctk\lib\scc* directories contain library files that are used to build SCC applications. The *msvcasm* subdirectories are used when building applications with MSVC++ and the *vacppasm* subdirectories are used when building applications with VACPP.

<sup>6</sup> In particular, the public keys and owner identifiers associated with segments 2 and 3 and the image names associated with segments 1 and 2 must have the values CFCrrrss.CLU and TDVrrrss.CLU assign them.

<sup>7</sup> In particular, the public key and owner identifier associated with segment 2 must have the values TDVrrrss.CLU assigns them.

- The *scctk\obj* directories are empty. The makefiles in the *scctk\src\* directories place object and executable files in this subtree.
- The *scctk\src\samples* directories contain the source for some sample host and SCC applications.
- The *scctk\src\udx* directory contains files used to create extensions to IBM's CCA application. This directory is not created unless the UDX Toolkit is installed.

The appropriate compiler options should be used when building an SCC application to ensure the directories listed below are searched for include files in the order shown:

1. *scctk\cpqenv\nt\vacppmsm* (if building with VACPP) or  
*scctk\cpqenv\nt\msvcasm* (if building with MSVC++)
2. *scctk\include\scc*
3. *scctk\include\common*

For example, the compiler options below might be specified to build an SCC application with VACPP:

```
/Ic:\scctk\cpqenv\nt\vacppmsm /Ic:\scctk\include\scc /Ic:\scctk\include\common
```

Similarly, the appropriate compiler options should be used when building a host application to ensure the directories listed below are searched for include files in the order shown:

1. *scctk\include\host*
2. *scctk\include\common*

For example, the compiler options below might be specified to build a host application with MSVC++:

```
/Ic:\scctk\include\host /Ic:\scctk\include\common
```

---

## Preparing the Development Platform

After the Developer's Toolkit (and the UDX Toolkit, if appropriate) and all prerequisites (see "Prerequisites" on page 1-2) have been installed, the developer must prepare the coprocessor for use as a development platform. The specific procedure depends on whether or not software has already been installed in the coprocessor and, if so, what software has been installed.

CLU's ST command can be used to determine what software, if any, is loaded in the coprocessor. For example:

```
TKNCLU CLU.LOG ST
```

An excerpt from a typical response to this command is as follows:

```

*** ROM Status; INIT: INITIALIZED
*** ROM Status; SEG2: RUNNABLE , OWNER2: 03
*** ROM Status; SEG3: RUNNABLE , OWNER3: 06
*** Page 1 Certified: YES
*** Segment 1 Image: CCA 1.2.2 SEGMENT-1 ...
*** Segment 1 Revision: 122
*** Segment 2 Image: CP/Q++ 1.22 ...
*** Segment 2 Revision: 122
*** Segment 3 Image: ...
*** Segment 3 Revision: 1

```

### The First ROM Status Line

If the first “ROM Status” line does not indicate segment 1 is in the INIT state or if page 1 is not certified, the coprocessor cannot be used as a development platform without additional assistance from IBM.

### Segments 2 and 3 UNOWNED

If the “ROM Status” line indicate segments 2 and 3 are UNOWNED, the contents of segment 1 (as specified in the “Segment 1 Image” line) dictate how to proceed:

- **Coprocessor in Factory-Default State** - If software has never been loaded into the coprocessor (for example, if the coprocessor has just been removed from a factory-sealed package), the segment 1 image name will likely be rather cryptic. In this case, the developer updates the coprocessor’s system software by loading CFCrrss.CLU into the coprocessor, for example:

```
TKNCLU CLU.LOG PL CFCrrss.CLU
```

If this command fails, further assistance from IBM is required. (The failure may indicate the public key associated with segment 1 has not been set to the expected factory default.)

If this command succeeds, the developer proceeds to load TDVrrss.CLU as indicated in “Segment 1 Current” below.

- **Segment 1 Downlevel** - If segment 1 contains a downlevel version or revision of CCA segment 1, the developer must first reset the coprocessor to its factory-default state by loading the copy of C2Frrss.CLU that corresponds to the version and revision of the CCA segment 1 loaded on the card<sup>8</sup>, for example:

```
TKNCLU CLU.LOG PL C2Fxxxyy.CLU
```

where xxx is the downlevel version of CCA and yy is the downlevel revision. The developer then proceeds to load CFCrrss.CLU as indicated in “Coprocessor in Factory-Default State” above.

- **Segment 1 Current** - If segment 1 contains the appropriate version and revision of CCA segment 1, the developer prepares the coprocessor for use as a development platform by loading TDVrrss.CLU into the coprocessor, for example:

```
TKNCLU CLU.LOG PL TDVrrss.CLU
```

<sup>8</sup> If the downlevel CCA segment 1 is version 1.22 or earlier, additional assistance from IBM is required to update the system software.



If desired, the developer can confirm the software has been properly loaded by resetting the coprocessor to start the “reverse-then-echo” application loaded by TDVrrrss.CLU (see Appendix C, “How to Reboot the IBM 4758” on page C-1 for details) and then running the host reverse-then-echo driver, for example:

**HRE *text***

The driver sends *text* to the reverse-then-echo application on the coprocessor, which reverses it and returns it to the driver. The driver prints the text received.

The developer then proceeds to load TPRrrrss.CLU as indicated in “Segment 2 Owner ID = 3 and Segment 3 Owner ID = 6” below.

### Segments 2 and 3 RUNNABLE

If the “ROM Status” lines indicate segments 2 and 3 are RUNNABLE, the owner identifiers specified on those lines for segments 2 and 3 dictate how to proceed:

- **Segment 2 Owner ID = 2** - If the owner identifier associated with segment 2 is 2, the developer relinquishes ownership of segment 2 by loading CRSrrrss.CLU into the coprocessor,<sup>9</sup> for example:

**TKNCLU CLU.LOG PL CRSrrrss.CLU**

If this command fails, further assistance from IBM is required. (The failure may indicate the public key associated with segment 2 has not been set to the expected value.)

If this command succeeds, segments 2 and 3 become UNOWNED and the developer proceeds according to the instructions given in “Segments 2 and 3 UNOWNED” above.<sup>10</sup>

- **Segment 2 Owner ID = 3 and Segment 3 Owner ID = 6** - If the owner identifier associated with segment 2 is 3 and the owner identifier associated with segment 3 is 6, the developer loads into the coprocessor TPRrrrss.CLU, which contains a version of CP/Q++ that supports the debugging of applications, for example:

**TKNCLU CLU.LOG PL TPRrrrss.CLU**

This completes preparation of the coprocessor for use as a development platform.

- **Other Owner IDs** - If the owner identifiers associated with segment 2 and/or 3 differ from those previously listed, it may not be possible to use the coprocessor for development. To do so requires the assistance of the owner of segment 2, who must supply a CLU file to surrender that ownership.

### Segments 2 and 3 Neither UNOWNED nor RUNNABLE

If the “ROM Status” lines indicate segments 2 and 3 are OWNED\_BUT\_UNRELIABLE or RELIABLE\_BUT\_UNRUNNABLE, the coprocessor cannot be used as a development platform without additional assistance from IBM.

<sup>9</sup> CRSrrrss.CLU ships with the IBM CCA Support Program.

<sup>10</sup> If CRS12200.CLU is used to relinquish ownership of segment 2, the developer must load CFCrrrss.CLU as indicated in “Segments 2 and 3 UNOWNED.” The contents of segment 1 cannot be used as a guide in this case.



---

## Chapter 3. Developing and Debugging an SCC Application

This chapter describes how to use the software in the development environment to create an SCC application on Windows NT and prepare it to be loaded into an IBM 4758 coprocessor and debugged.

This chapter describes:

- Each step in the development process
- Special coding requirements for development
- Unsupported CP/Q base operating system function calls
- Supported and unsupported C run-time library function calls and global variables
- Required option and switch settings for the compiler, assembler, linker, and librarian
- How to convert a compiled SCC application into a version that CP/Q++ can load and execute
- How to build a read-only disk image containing the SCC application
- How to load the disk image into the coprocessor
- How to start the debugger

---

### Development Process Road Map

As introduced in Chapter 1, "Introduction," the procedure to build an SCC application and load it into the development coprocessor consists of the following steps:

1. Compile, assemble, and link
2. Translate
3. Build disk image
4. Load image into the coprocessor

Figure 3-1 on page 3-2 illustrates the development process, and indicates the name of the tool and input needed to perform each step. The process is identical to that shown in Figure 1-1 on page 1-3; this flowchart simply provides more detail.

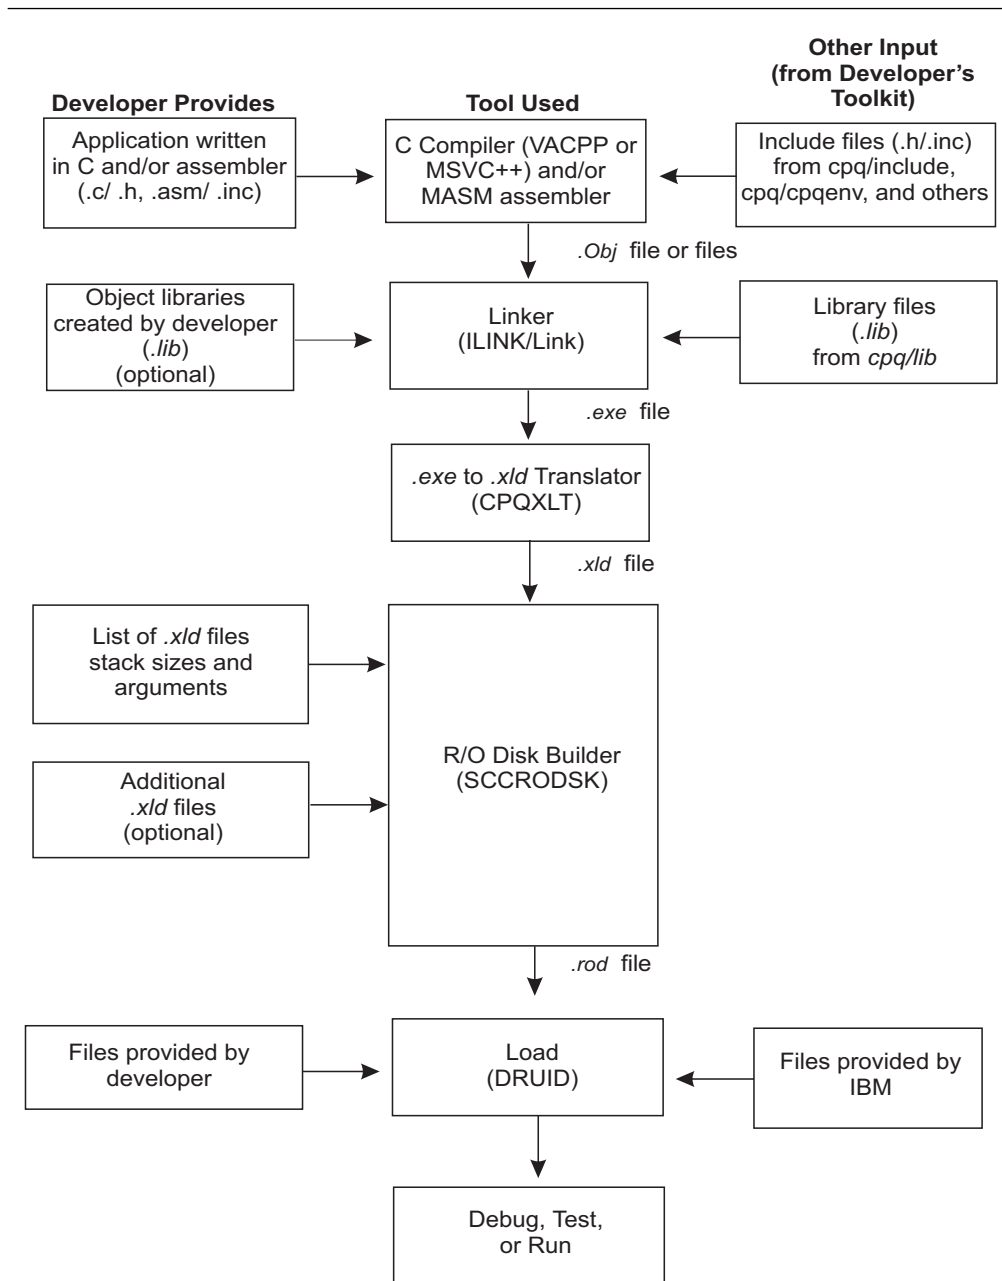


Figure 3-1. Development Process Road Map

The following sections detail how to use the Developer's Toolkit to perform these steps.

---

## Special Coding Requirements During Development

### Developer Identifiers

An SCC application must register with a CP/Q++ device manager before the application can receive requests from the host. The application must supply a “developer identifier” that uniquely identifies the developer as part of the registration process.<sup>1</sup> During development, a developer may use an arbitrary nonzero value for the developer identifier. Before an application can be released, the developer must obtain a unique identifier from IBM and must rebuild the application and any host application that interacts with it to use the true identifier.

### Attaching with the Debugger

An application that has been downloaded to the coprocessor will be loaded and start to run as soon as the coprocessor is rebooted and may run for some time before the debugger places the application under debug and quiesces it. To ensure the application does not make too much progress before the debugger takes control, the developer must code an infinite loop early in the application and use the debugger to move the execution point past the loop after the application is quiesced. To ensure the loop does not starve other agents in the system, the loop should be coded along the following lines:

```
QMSGHDR msg;
unsigned long count;

i = 0;
memset (&msg,0,sizeof(msg));
for (;;)
{
    CPRecvMsg(&msg,&count,0,1000000);
    i++;
}
```

The timeout in the call to CPRecvMsg should be large enough to allow other agents in the system to run most of the time but not so large that the call seldom returns.

After attaching to the application with the debugger, set a breakpoint on the `i++` statement and allow the application to run. When the breakpoint is hit, use the debugger's **Jump to location** function to move the execution point out of the loop.

---

## Compiling, Assembling, and Linking

The commercial compiler used in development of an SCC application is designed to create applications to run on a workstation under Windows NT rather than on a cryptographic coprocessor under CP/Q++. Consequently, the include files and libraries shipped with the compiler and the defaults for several options are not always appropriate for SCC applications.

---

<sup>1</sup> Refer to the description of `sccSignOn` in *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System Application Programming Reference* for details.

This section lists the base operating system and C run-time library function calls that an SCC application may use and those that are not supported.<sup>2</sup> It also lists options that must be specified when compiling, assembling, or linking to ensure that an SCC application will run properly. Other options may also be specified as long as they do not conflict with the options listed in this section.

The Developer's Toolkit includes makefiles that specify the proper options for each tool. Refer to *cpqenv.mak* in `lsccck\cpqenv\nt\vacppmsm` (when compiling with VACPP) or in `lsccck\cpqenv\nt\msvcmasm` (when compiling with MSVC) for details on their use.

## CP/Q Base Operating System Function Support

The Developer's Toolkit supports most of the functions described in *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System Application Programming Reference*. Exceptions are noted in Table 3-1. Furthermore, an SCC application has no special privileges and consequently cannot invoke certain restricted functions (for example, CPPTrace).

Table 3-1. *Unsupported CP/Q Functions*

CPSigReturn	CPSigVec
CPSigSend	
CPSigStack	
CPSigMask	
CPSigInt	
CPSigPreempt	

## C Run-Time Library Support

The Developer's Toolkit supports most of the library functions and global variables described in the *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System C Runtime Library Reference*.<sup>3</sup> However, several functions have been modified and others are not supported at all. Furthermore, there are restrictions on the use of certain intrinsic functions. See "Compiler Options" on page 3-6 for details.

### Supported Functions and Global Variables

Table 3-2 on page 3-5 lists the C run-time functions and global variables the Developer's Toolkit supports. Numbers after names refer to the notes following the table.

<sup>2</sup> Most unsupported functions are declared in an include file but are not implemented. That is, a program that invokes an unsupported function may compile, but it will not link.

<sup>3</sup> The C run-time library shipped with the Developer's Toolkit includes the `cinitnon` and `exitnon` versions of the initialization and exit code.

Table 3-2. Supported C Run-time Functions and Global Variables

_cdapage	abs	ffs	issupv	qsort (1)	strncmp
_exit	asctime	free	isupper	rand (5)	strncpy
_fullname	atexit	getenv	isxdigit	realloc	strnicmp
_isalnum	atof	getenvall	itoa	rindex	strpbrk
_isalpha	atoi	getenvall2	labs	setjmp	strrchr
_isascii	atol	getopt	ldiv	sprintf	strrev
_iscntrl	bcmp	gmtime (2)	localtime (3)	srand (5)	strspn
_isdigit	bcopy	hsort (1)	longjmp	sscanf	strstr
_isgraph	binsort (1)	index	malloc	strcat	strtok
_islower	bsearch (1)	insort (1)	memccpy	strchr	strtol
_isprint	bzero	isalnum	memchr	strcmp	strtoul
_ispunct	calloc	isalpha	memcmp	strcpy	strupr
_isspace	cjlvsn	isascii	memcpy	strcspn	swab
_isupper	clock	iscntrl	memicmp	strdup	time (6)
_isxdigit	copyenv	isdigit	memmove	strerror	tolower
_mons	ctime	isgraph	memset	strftime	toupper
_tolower	difftime	islower	mktime	stricmp	va_arg
_toupper	div	isprint	msort (1)	strlen	va_end
_wdays	errno	ispunct	printf (4)	strlwr	va_start
abort	exit	isspace	putenv	strncat	vsprintf

### Notes

1. This function takes as an argument the address of a comparison function. The comparison function must use the `__cdecl` linkage convention.
2. `gmtime` always sets the `tm_isdst` field of the output structure to 0.
3. `localtime` assumes local standard time is 300 minutes behind (west) of Greenwich Mean Time and makes an appropriate adjustment for daylight savings time.
4. The debug version of CP/Q++ forwards output generated by `printf` to the debugger, which displays the output in the Messages Window. The production version of CP/Q++ treats `printf` as a (relatively expensive) NOP.
5. `rand` and `srand` are ANSI-standard pseudo-random functions. Developers writing SCC applications should use `sccGetRandomNumber` instead.
6. `time` assumes the coprocessor time-of-day clock is set to local time (i.e., local standard time with an appropriate adjustment for daylight savings time) and that local standard time is 300 minutes behind (west) of Greenwich Mean Time.

### Unsupported Functions and Global Variables

Table 3-3 lists the C run-time functions and global variables the Developer's Toolkit does not support. Numbers after names refer to the notes following the table.

Table 3-3. Unsupported C Functions and Global Variables

_job	fflush	freopen	getsessid	read	signal
assert	fgetc	fscanf	lseek	remove	system
brkpt	fgetpos	fseek	open	rename	tell
close	fgets	fsetpos	paws	rewind	tmpfile
clearerr (1)	fopen	ftell	perror	scanf	tmpnam
feof (1)	fprintf	fwrite	putc (3)	setattr	ungetc
ferror (1)	fputc	getc (2)	putchar	setbuf	vfprintf
fileno (1)	fputs	getchar	puts	setsessid	vprintf
fclose	fread	gets	raise	setvbuf	write

**Notes**

1. This function is implemented as a macro and consequently programs that invoke it will compile and link. However, the C run-time library does not define any FILE structures that the function could take as an argument.
2. Use of `getc` will cause the link to fail with `getctext` and `__fillbuf` undefined.
3. Use of `putc` will cause the link to fail with `putctext` and `__flushbu` undefined.

**Compiler Options**

The Developer's Toolkit supports two compilers: IBM VisualAge C++ (VACPP) and Microsoft Visual C++ (MSVC++).<sup>4</sup>

**Reminder:** Although you use a C++ compiler to compile SCC applications, the applications must be written in C.

**VisualAge C++ (VACPP) Options**

When using VACPP to compile executable files, use the following switches with the `icc` command to control the process:

<b>Switch</b>	<b>Function</b>
<code>/Gn+</code>	Do not put default library information into the object file.
<code>/Gs+</code>	Do not generate stack probes.
<code>/I</code>	Use Developer's Toolkit customized settings to search for include files when compiling. See "Directories and Files" on page 2-1 for a description of the default file locations.
<code>/O</code>	Compile without optimization ( <code>/O-</code> ) when creating an executable suitable for debugging. Optimization may be enabled ( <code>/O+</code> ) when creating a production version.
<code>/Rn</code>	Do not incorporate the VACPP run-time environment into the compiled module.
<code>/Ti</code>	Generate debug information ( <code>/Ti+</code> ) when creating an executable suitable for debugging. Debug information may be omitted ( <code>/Ti-</code> ) when creating a production version.
<code>/Xi</code>	Do not search the path specified by the <code>include</code> environment variable.

**Note:** Developers writing extensions for IBM's CCA application must ensure CCA services are invoked using the `__stdcall` calling convention. One way to accomplish this is to specify the `/Mt` option with the `icc` command.

**Notes on Intrinsic Functions:** The Developer's Toolkit supports most VACPP intrinsic (automatically inlined) functions; the only unsupported intrinsic function is `__getTIBvalue`.

The following string and memory functions have intrinsic forms:

---

<sup>4</sup> Developers writing extensions for IBM's CCA application must use IBM VisualAge C++.



Table 3-4. VACPP String and Memory Functions with Intrinsic Forms

memcpy	strcat	strncat
memchr	strchr	strncmp
memcmp	strcmp	strncpy
memset	strcpy	strchr
memmove	strlen	

To use the intrinsic forms of these functions, compile with optimization (/O+) and define `_STRING_INTRINSICS_` before including `<string.h>`; define `_STRING_INTRINSICS_` in your code with the following syntax:

```
#define _STRING_INTRINSICS_
```

As an alternative, define `_STRING_INTRINSICS_` when compiling by using the `/D_STRING_INTRINSICS_` compiler option.

**Note:** Do not define `_STRING_INTRINSICS_` unless optimization is enabled.

The nonintrinsic forms of the following functions are not supported:

Table 3-5. VACPP Functions Whose Nonintrinsic Form Is Not Supported

<code>_clear87</code>	<code>_control87</code>	<code>_status87</code>
-----------------------	-------------------------	------------------------

The following functions are not supported:

Table 3-6. Unsupported VACPP Functions

wscat	wscopy	wscncmp
wchr	wcslen	wscncpy
wscmp	wscncat	wscrchr

**Note:** Routines that invoke `setjmp` must be compiled without optimization.

### Microsoft Visual C++ (MSVC++) Options

When using MSVC++ to compile executable files, use the following switches with the `cl` command to control the process:

Switch	Function
<code>/Gs100000000</code>	Do not generate stack probes. While the value entered does not need to be '100000000', it does need to exceed the size of any stack frame possible.
<code>/I</code>	Use Developer's Toolkit customized settings to search for include files when compiling. See "Directories and Files" on page 2-1 for a description of the default file locations.
<code>/O</code>	Compile without optimization (/Od) when creating an executable suitable for debugging. Optimization may be enabled (/Ox) when creating a production version.
<code>/X</code>	Do not search the path specified by the <code>include</code> environment variable.
<code>/Z7</code>	Generate debug information (/Z7) when creating an executable suitable for debugging. Debug information must be incorporated into the executable, not placed in a program database (that is, do not specify /Zi). Debug information may be omitted when creating a production version.

`/ZI` Do not put default library information into the object file.

**Notes on Intrinsic Functions:** The following MSVC++ intrinsic functions are always generated inline:

Table 3-7. MSVC++ Intrinsic Functions (Always Inlined)

<code>_disable</code>	<code>_lrotl</code>	<code>_rotl</code>
<code>_enable</code>	<code>_lrotr</code>	<code>_rotr</code>
<code>_inp</code>	<code>_outp</code>	<code>_strset</code>
<code>_inpd</code>	<code>_outpd</code>	
<code>_inpw</code>	<code>_outpw</code>	

Code the appropriate intrinsic compiler directive `#pragma` or enable optimization to use the inline forms of the following intrinsic functions:

Table 3-8. MSVC++ Intrinsic Functions (Inlined by Directive or Optimization)

<code>abs</code>	<code>memcpy</code>	<code>strcpy</code>
<code>fabs</code>	<code>memset</code>	<code>strlen</code>
<code>labs</code>	<code>strcat</code>	
<code>memcmp</code>	<code>strcmp</code>	

The following intrinsic functions are inlined if the appropriate optimization level is in effect:

Table 3-9. MSVC++ Intrinsic Functions (Inlined by Optimization)

<code>atan</code>	<code>exp</code>	<code>sin</code>
<code>atan2</code>	<code>log</code>	<code>sqrt</code>
<code>cos</code>	<code>log10</code>	<code>tan</code>

**Notes:**

1. While each function listed in Table 3-9 has a non-intrinsic form, the Developer's Toolkit supports only the non-intrinsic forms of **log** and **log10**.
2. The `_alloca` function is not supported.
3. Routines that invoke `setjmp` must be compiled without optimization.

## Assembler Options

The Developer's Toolkit supports the Microsoft\*\* MASM assembler. Use the following switches with the `m1` command to control the process:

Switch	Function
<code>/coff</code>	Generate output as Common Object File Format (COFF), rather than Object Module Format (OMF). This switch should only be used when using the Microsoft Visual C++ compiler and the LINK linker.
<code>/Cp</code>	Preserve case of identifiers.
<code>/I</code>	Set the include path. See "Directories and Files" on page 2-1 for a discussion of how to set these options.
<code>/X</code>	Do not search the path specified by the <code>include</code> environment variable.
<code>/Zd /Zi</code>	Generate debug information ( <code>/Zd</code> and <code>/Zi</code> ) when creating an executable suitable for debugging. <sup>5</sup> Debug information may be omitted when creating a production version.

## Linker Options

The compiler used determines which linker must be used to create an executable file from the resulting object (*.obj*) files.

### ILINK (VACPP Linker)

ILINK links object files created by the VACPP compiler. Use the switch below with the *ilink* command to control the process:

Switch	Function
<i>/debug</i>	Preserve debug information in the object files (/DEBUG) when creating an executable suitable for debugging. Debug information may be omitted (/NODEBUG) when creating a production version.
<i>/nod</i>	Do not search default libraries listed in the object files.
<i>/noe</i>	Do not search library extended directories.
<i>/pm:vio</i>	Specify a .EXE type. The type specified (PM, VIO, or NOVIO) does not matter, but if this switch is not provided the linker issues a warning message.

The libraries supplied in *scctl\lib\scclnt\vacppasm* were compiled with the /GI+ option. Use the /OPTFUNC linker option to remove unreferenced functions and generate a smaller executable file.

### LINK (MSVC++ Linker)

LINK links object files created by the MSVC++ compiler. Use the switches below with the *link* command to control the process:

Switch	Function
<i>/debug</i>	Preserve debug information in the object files (/DEBUG) when creating an executable suitable for debugging. If this option is specified, /DEBUGTYPE:CV and /PDB:none must also be specified. Debug information may be omitted when creating a production version.
<i>/entry:startup</i>	Specify the program entry point.
<i>/fixed:no</i>	Preserve relocation information in the executable file.
<i>/nodefaultlib</i>	Do not search the default libraries.
<i>/stack:0x1000000</i>	Set the initial stack size.

The libraries supplied in *scctl\lib\scclnt\msvcasm* were compiled with the /Gy option. Use the /OPT:REF linker option to remove unreferenced functions and generate a smaller executable file.

<sup>5</sup> Debug information generated by MASM is not compatible with the format expected by the linker that ships with VisualAge C++. Do not use the /Zd or /Zi switches when assembling files that will be linked using that linker.

## Librarian Options

The compiler type used determines which librarian must be used to create a library (.lib) file from one or more object (.obj) files. ILIB creates libraries from files generated by VACPP, and LIB creates libraries from files generated by MSVC++.

There are no required option settings for either librarian.

---

## Translating

The Translator Utility (CPQXLT.EXE) translates a fully-compiled executable file into executable file able to run on the CP/Q++ operating system embedded within the coprocessor. The utility supports pure 32-bit executable files built from C or assembler source code; it does not support dynamic link libraries or C++ programs.

Debug information is translated for Windows NT executables built with the compilers described above, and for executables built with Microsoft-compatible debug information. The compiler, assembler, and linker must have been invoked with the proper options in order to generate debug information (for example, VACPP /Ti+ or MSVC++ /Z7).

Because the entire executable (including any debug information, if present) is incorporated into the read-only disk image that is loaded into the coprocessor, it is recommended that the translation be performed twice, once to create a version of the executable containing debug information for the debugger's use, and a second time to create a version without debug information to be downloaded to the coprocessor.

### Syntax

```
cpqxlt [input-filename] [output-filename] {optional switches}
```

The optional switches are as follows:

Switch	Function
<i>/base:address</i>	Sets the address for the first loaded section. (The default is 0).
<i>/nodebug</i>	Suppresses translation of debug information.
<i>/align:factor</i>	Sets the alignment boundary for loaded sections. (The default is a 512-byte boundary).

---

## Building Read-Only Disk Images

The Disk Builder Utility (SCCRODSK.EXE) creates a read-only disk image that can be loaded into the coprocessor using DRUID or can be signed using TKNSGMR and placed into a CLU file by TKNPKGR for subsequent download by CLU.

To build a disk image, create an ASCII text file (*inputfile.txt*) that lists the files to incorporate into the disk image.<sup>6</sup> Each line in the file has the form:

---

<sup>6</sup> At present, only one .xld file can be loaded into the coprocessor. Any others listed are incorporated into the disk image but are not loaded.

```
filename [stacksize] [arg1 [arg2 [...]]]
```

where *filename* is the name (or full pathname if desired) of the file to incorporate into the disk image, *stacksize* is the number of bytes to allocate for the stack when the application the file contains is loaded and run. (The default is 4096 bytes.) Any additional tokens on the line (*arg1*, *arg2*, and so on) are passed to the application as invocation arguments. Blank lines are ignored and lines containing an asterisk in the first column are treated as comments.

For example, the following line in the input file

```
rte.xld 8192 a b cde "space " 'quote' " " 1 2 3 4
```

causes *rte.xld* to be run with an 8K stack. On entry to `main ( )`, `argc` and `argv` have the following values:

```
argc      11
argv[0]   "rte.xld"
argv[1]   "a"
argv[2]   "b"
argv[3]   "cde"
argv[4]   "space"
argv[5]   "'quote'"
argv[6]   " "
argv[7]   "1"
argv[8]   "2"
argv[9]   "3"
argv[10]  "4"
```

After the ASCII text file has been created, invoke the disk builder utility as follows:

```
sccrodsk inputfile.txt outputfile.rod
```

The utility creates the disk image.

---

## Downloading and Debugging

Once a file containing the read-only disk image of the application has been generated by SCCRODSK, the file may be downloaded to the coprocessor using DRUID.

DRUID does not affect any data in the nonvolatile memory (battery-backed RAM and flash) associated with the application. If the developer wishes to clear state that has accumulated during prior debug sessions so that the application will start with a clean slate, the developer should first download TL3rrrs.CLU to the coprocessor using CLU:

```
TKNCLU CLU.LOG PL TL3rrrs.CLU
```

### Syntax

```
druid [image_fn pubkey_fn privkey_fn [coprocessor_number]]
```

where

- *image\_fn* is the name of the file containing the read-only disk image to download to the coprocessor.

- *pubkey\_fn* is the name of a file containing the public key to be associated with the application<sup>7</sup> (for example, S3KCLRPU.DRK).
- *privkey\_fn* is the name of a file containing an RSA keypair (for example, S3KCLRPP.DRK). The public key must match the public key currently associated with the application in the coprocessor.<sup>8</sup>
- More than one coprocessor may be installed in a host. *coprocessor\_number* identifies the coprocessor to which the read-only disk image is downloaded. The default is 0.

The number assigned to a particular coprocessor depends on the order in which information about devices in the system is presented to the device driver by the host operating system. At the present time there is no way to tell *a priori* which coprocessor will be assigned a given number.

If DRUID is invoked without arguments, it prompts for them.

DRUID displays a summary of the status of the coprocessor before it downloads the application. The summary includes

- The coprocessor's serial number<sup>9</sup>
- The current left and right bootcounts (see "Targeting Arguments" on page E-17 for details)
- The name, creation date, and size of the image file last downloaded to the coprocessor
- The name of the file containing the public key associated with the application currently loaded in the coprocessor<sup>10</sup>

CP/Q++ will load and run the application after the coprocessor is rebooted. See Appendix C, "How to Reboot the IBM 4758" on page C-1 for a description of how to reboot the coprocessor.

Details on the use of CLU can be found in the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Installation Manual*.

After the application is running, it can be debugged using the ICAT debugger. Refer to the *IBM 4758 PCI Cryptographic Coprocessor Interactive Code Analysis Tool (ICAT) User's Guide* for details.

---

<sup>7</sup> That is, the public key to be associated with segment 3.

<sup>8</sup> That is, the public key currently associated with segment 3.

<sup>9</sup> That is, the value `sccGetConfig` returns in `pInfo->VPD.AdapterID`. Refer to the *IBM 4758 PCI Cryptographic Coprocessor CP/Q Operating System Application Programming Reference* for details.

<sup>10</sup> That is, the value of *pubkey\_fn* supplied when DRUID last downloaded an application to the coprocessor.

---

## Chapter 4. Testing an SCC Application in a Production Environment

The version of CP/Q++ that allows an application to be debugged includes certain components that are not present in the production version of CP/Q++. A developer may find it prudent to test the application running under a production version of CP/Q++ before releasing the application. The TNPrrss.CLU file shipped with the Developer's Toolkit can be used to replace the debug version of CP/Q++ with a production version:

```
TKNCLU CLU.LOG PL TNPrrss.CLU
```

The developer should then rebuild the application (*without debug information or any code added for debugging purposes, including any infinite loop added to allow the debugger to attach*), create a read-only disk image, and download it to the coprocessor in the same manner as described in chapter 3.

If further debugging proves necessary, the debug version of CP/Q++ can be reloaded into the coprocessor as follows:

```
TKNCLU CLU.LOG PL TPRrrss.CLU
```

Use of TNPrrss.CLU and TPRrrss.CLU in this manner preserves any state information the application has saved in battery-backed RAM and flash. "Downloading and Debugging" on page 3-11 describes how to clear such state information (if desired) before downloading an application.





---

## Chapter 5. Packaging and Releasing an SCC Application

The design for the IBM 4758 PCI Cryptographic Coprocessor was motivated by the need to simultaneously satisfy the following requirements<sup>1</sup>:

1. Code must not be loaded into the coprocessor unless IBM or an agent IBM trusts has authorized the operation.
2. Once loaded into the coprocessor, code must not run or accumulate state unless the environment in which it runs is trustworthy.
3. Agents outside the coprocessor that interact with code running on the coprocessor must be able to verify that the code is legitimate and that the coprocessor is authentic and has not been tampered with.
4. Shipment and configuration of coprocessors and maintenance on and upgrades to code inside a coprocessor must not require trusted couriers or security officers.
5. IBM must not need to examine a developer's code or have any knowledge of a developer's private cryptographic keys in order to make it possible for customers to load the developer's code into a coprocessor and run it.

To meet these requirements, the design defines four "segments":

- Segment 0 is ROM and contains one portion of "Miniboot". Miniboot is the most privileged software in the coprocessor and among other things implements the security protocols described in this section.
- Segment 1 is flash and contains the other portion of "Miniboot". The division of Miniboot into a ROM portion and a flash portion preserves flexibility (the flash portion can be changed if necessary) while guaranteeing a basic level of security (implemented in the ROM portion).
- Segment 2 is flash and usually contains the coprocessor operating system.
- Segment 3 is flash and usually contains one or more coprocessor applications.

The security protocols that enforce these design goals are based on RSA keypairs and a notion of who owns the code in each segment. IBM owns segments 1 and 2 and issues an owner identifier to any party that is developing code to be loaded into segment 3. The coprocessor saves the identity of the owner of each segment and an RSA public key for each segment. The key is provided by the segment's owner.

The coprocessor will not accept a command that changes the contents of a segment unless the command is digitally signed with the private key that corresponds to the public key associated with the segment. The command must also correctly identify the owner of the segment. Commands that must change the contents of a segment that does not yet have a public key must be signed with the private key that corresponds to the public key associated with the segment's parent. For example, the command that initially sets the contents, owner, and public key for segment 3 must be signed with the private key for segment 2.

---

<sup>1</sup> For a thorough overview of the coprocessor's security goals and a description of the security architecture, refer to *Building a High-Performance, Programmable Secure Coprocessor*, Research Report RC21102 published by the IBM T.J. Watson Research Center in February 1998.

The files shipped in the Developer's Toolkit are designed to make it very easy for a developer to start work immediately but are also constructed in a way that does not threaten the security or integrity of an application deployed in the field or one that may be deployed in the future. During development, the developer uses a default RSA keypair (which makes development easy) that is tied to a generic owner identifier (which makes the generic keypair "harmless"). When the developer is ready to deploy an application in the field, the developer must obtain a unique developer identifier from IBM and must generate a new, unique RSA keypair. This is summarized in the table below.

Attribute	Development	Production
Owner	"Generic developer"	Developer-unique identifier
Public Key	Generic (common) key	Developer-generated key

Prior to deployment, a developer must restore the coprocessor used for development to a state suitable for use in production<sup>2</sup> using TRSrrrs.CLU:

**TKNCLU CLU.LOG PL TRSrrrs.CLU**

The developer must then install the CCA Support Program on the host, install the CCA application on the coprocessor, and configure a CCA test node. Instructions on how to complete these steps appear in chapters 3, 4, and 5, respectively, of the *IBM 4758 PCI Cryptographic Coprocessor CCA Support Program*. This prepares the coprocessor for use by TKNSGMR and TKNPKGR.

The developer generates three RSA keypairs using TKNSGMR's KEYGEN function, for example:

**TKNSGMR KEYGEN 2 S3KDEVPP.KEY S3KDEVPU.KEY DFT\_SKEL.TKN  
TKNSGMR KEYGEN 2 DEVSGNPP.KEY DEVSGNPU.KEY DFT\_SKEL.TKN  
TKNSGMR KEYGEN 2 DEVPKGPP.KEY DEVPKGPU.KEY DFT\_SKEL.TKN**

The first keypair supplies the key to be saved with the developer's application in segment 3. The second and third keypairs are used by TKNSGMR and TKNPKGR, respectively, to generate digital signatures that CLU uses to verify that IBM has authorized its use.

The KEYGEN function creates two KEY files, one containing both the private and public keys (for example, S3KDEVPP.KEY) and the other containing just the public key (for example, S3KDEVPU.KEY). The KEYGEN function also creates a file containing the hash of the public key. The file has the same name as the file containing the public key and an extension of HSH (for example, S3KDEVPU.HSH). The developer forwards each public key file to IBM. The developer should also communicate the hash value of each public key to IBM (by way of a separate channel) to ensure an adversary has not replaced the developer's public key file with another.

After an appropriate contract has been signed, IBM supplies a file (DEV3.TXT) containing a unique owner identifier for the developer and a file (IBMCLT2.TXT) containing the owner identifier associated with segment 2. IBM also creates

<sup>2</sup> Its factory default state or obtain a second IBM 4758.

- an emergency signature file (ESIGDEV.SIG) incorporating the developer's owner identifier and segment public key and
- a TKNSGMR input file (DEVrrrss.TSK) that loads CP/Q++ into an IBM 4758 shipped from the factory and sets the owner identifier associated with segment 3 to the developer's owner identifier.

Finally, IBM supplies certificates for the TKNSGMR and TKNPGR public files (DEVSGNPU.CRT and DEVPKGGPU.CRT, respectively).

IBM generates a hash of each file (DEV3.HSH, IBMCLT2.HSH, ESIGDEV.HSH, DEVrrrss.HSH, DEVSGNPU.HSH, and DEVPKGGPU.HSH) and forwards the files and the hashes (by way of a separate channel) to the developer, who should verify that the hashes match the files they cover using TKNSGMR's HASH\_VER function, for example:

```
TKNSGMR HASH_VER DEV3.HSH DEV3.TXT
TKNSGMR HASH_VER IBMCLT2.HSH IBMCLT2.TXT
TKNSGMR HASH_VER ESIGDEV.HSH ESIGDEV.SIG
TKNSGMR HASH_VER DEVrrrss.HSH DEVrrrss.TSK
TKNSGMR HASH_VER DEVSGNPU.HSH DEVSGNPU.CRT
TKNSGMR HASH_VER DEVPKGGPU.HSH DEVPKGGPU.CRT
```

The developer then builds a version of the application for release (for example, builds without debugging information or debug code and changes the value of pAgentID->DeveloperID in any calls to sccSignOn and the value of pRequestBlock->AgentID.DeveloperID in any calls to sccRequest to the number supplied in DEV3.TXT) and uses TKNSGMR to create an EMBURN3 command that incorporates the application, IBM's segment 2 owner ID, the developer's owner ID, and the developer's unique keys, for example:

```
TKNSGMR EMBURN3 MYAPP.TSK
  part version description
  DEVSGNPU.CRT DEVSGNPP.KEY
  APP.ROD title revision
  S3KDEVPP.KEY ESIGDEV.SIG
  ibmclt2 dev3
  1 1
  a 0 b 0 c 0 d 0 e 0 0
  x 0 0 65535 0 0
  x 0 0 65535 0 0
```

where *part*, *version*, and *description* supply information that is incorporated into the output file, *title* and *revision* supply information that is downloaded to the coprocessor and stored with the application in segment 3, *ibmclt2* is the number supplied in IBMCLT2.TXT, and *dev3* is the number supplied in DEV3.TXT. See Appendix E, "Using Signer and Packager" on page E-1 for details.

The developer uses TKNPGR to combine the task file IBM supplies with the file containing the EMBURN3 command, for example:

```
TKNPGR DEVPKGGPU.CRT DEVPKGPP.KEY
  2 DEVrrrss.TSK MYAPP.TSK MYAPP.CLU
  part version description
```

where *part*, *version*, and *description* supply information that is incorporated into the output file. See Appendix E, "Using Signer and Packager" on page E-1 for details.

The file generated by TKNPKGR can be shipped to end users, who can use CLU to load the application into the coprocessor and start execution, for example:

```
TKNCLU CLU.LOG PL MYAPP.CLU
```

```
TKNCLU CLU.LOG RS
```

---

## Appendix A. An Overview of the Development Process

This appendix describes the entire process from initial preparation of the coprocessor to the creation of a file containing a developer application that can be shipped to the developer's customers or end users.

Each step in this overview is listed under a heading that notes where in the body of the manual the step or tools it uses is described.

### Preparing the Development Platform

1. Determine whether or not the coprocessor is empty:

#### **TKNCLU CLU.LOG ST**

If coprocessor segment 1 is not in the INIT state or if page 1 is not certified, the coprocessor cannot be used as a development platform without additional assistance from IBM.

If coprocessor segment 2 is UNOWNED, continue with step 2.

If coprocessor segment 2 is RUNNABLE and the owner identifier associated with segment 2 is 2, continue with step 3.

If coprocessor segment 2 is RUNNABLE and the owner identifier associated with segment 2 is 3, continue with step 4.

If coprocessor segment 2 is RUNNABLE but the owner identifier associated with segment 2 is neither 2 nor 3, it may be possible to use the coprocessor for development. To do so requires the assistance of the owner of segment 2, who must supply a CLU file to surrender that ownership.

2. If coprocessor segment 2 is UNOWNED, the contents of segment 1 dictate how to proceed:

- **Coprocessor in Factory-Default State** - If software has never been loaded into the coprocessor (for example, if the coprocessor has just been removed from a factory-sealed package), the segment 1 image name will likely be rather cryptic. In this case, update the coprocessor's system software by loading CFCrrrss.CLU into the coprocessor, for example:

#### **TKNCLU CLU.LOG PL CFCrrrss.CLU**

If this command fails, further assistance from IBM is required. (The failure may indicate the public key associated with segment 1 has not been set to the expected factory default.)

If this command succeeds, load TDVrrrss.CLU as indicated in "Segment 1 Current" on page A-2.

- **Segment 1 Downlevel** - If segment 1 contains a downlevel version or revision of CCA segment 1, reset the coprocessor to its factory-default state by loading the copy of C2Frrrss.CLU that corresponds to the version and revision of the CCA segment 1 loaded on the card<sup>1</sup>, for example:

#### **TKNCLU CLU.LOG PL C2Fxxxxyy.CLU**

---

<sup>1</sup> If the downlevel CCA segment 1 is version 1.22 or earlier, additional assistance from IBM is required to update the system software.

where *xxx* is the downlevel version of CCA and *yy* is the downlevel revision. Then load CFCrrss.CLU as indicated in “Coprocessor in Factory-Default State” on page A-1.

- **Segment 1 Current** - If segment 1 contains the appropriate version and revision of CCA segment 1, prepare the coprocessor for use as a development platform by loading TDVrrss.CLU into the coprocessor, for example:

**TKNCLU CLU.LOG PL TDVrrss.CLU**

If desired, confirm the software has been properly loaded by resetting the coprocessor to start the “reverse-then-echo” application loaded by TDVrrss.CLU (see Appendix C, “How to Reboot the IBM 4758” on page C-1 for details) and then running the host reverse-then-echo, for example:

**HRE *text***

The driver sends *text* to the reverse-then-echo application on the coprocessor, which reverses it and returns it to the driver. The driver prints the text received.

Continue with step 4.

3. If the owner identifier associated with coprocessor segment 2 is 2, relinquish ownership of segment 2 by loading CRSrrss.CLU into the coprocessor, for example:

**TKNCLU CLU.LOG PL CRSrrss.CLU**

If this command fails, further assistance from IBM is required. (The failure may indicate the public key associated with segment 2 has not been set to the expected value.)

If this command succeeds, segments 2 and 3 become UNOWNED. Continue with step 2.<sup>2</sup>

4. If the owner identifier associated with coprocessor segment 2 is 3 and the owner identifier associated with segment 3 is 6, load into the coprocessor TPRrrss.CLU, which contains a version of CP/Q++ that supports the debugging of applications, for example:

**TKNCLU CLU.LOG PL TPRrrss.CLU**

This completes preparation of the coprocessor for use as a development platform. Continue with step 5.

If the owner identifier associated with segment 2 is 3 but the owner identifier associated with segment 3 is not 6, relinquish ownership of segment 2 by loading TRSrrss.CLU into the coprocessor, for example:

**TKNCLU CLU.LOG PL TPRrrss.CLU**

Continue with step 2.

## Compiling, Assembling, and Linking

---

<sup>2</sup> If CRS12200.CLU is used to relinquish ownership of segment 2, the developer must load CFCrrss.CLU as indicated in step 2. The contents of segment 1 cannot be used as a guide in this case.

5. Compile and link the application under development. Specify the appropriate options to ensure debugging information is incorporated into the .EXE file produced (APP.EXE<sup>3</sup>).

### Translating

6. Translate the application to the CP/Q++ executable format:<sup>4</sup>

**CPQXLT APP.EXE APP.XLD**

Translate the application a second time, omitting any debug information:<sup>4</sup>

**CPQXLT APP.EXE DOWNLOAD\APP.XLD /NODEBUG**

The output files generated in the two translations must have the same name (for example, APP.XLD) and so must be placed in separate directories.

### Building Read-Only Disk Images

7. Create or modify the text file (RODISKIN.TXT<sup>5</sup>) that lists the names of the executable files to be loaded and run. In our example, if the developer wanted an 8K stack, the contents of RODISKIN.TXT might be:

```
c:\scctk\obj\app\nt\msvcasm\download\app.xld 8192
```

8. Build a read-only disk image that incorporates the application:<sup>6</sup>

**SCCRODSK RODISKIN.TXT APP.ROD**

### Downloading and Debugging

9. If desired, clear any state the application saved in nonvolatile memory during previous debug sessions:

**TKNCLU CLU.LOG PL TL3rrrss.CLU**

10. Download the file generated in step 8 to the coprocessor:<sup>7</sup>

**DRUID APP.ROD C:\SCCTK\ETC\S3KCLRPU.DRK C:\SCCTK\ETC\S3KCLRPP.DRK**

11. Wait for the coprocessor to reboot and start the application.

12. Start the debugger and attach to the application:

**ICATCPW**

Refer to the *IBM 4758 PCI Cryptographic Coprocessor Interactive Code Analysis Tool (ICAT) User's Guide* for more information.

If changes to the application prove necessary, make them and continue with step 5.

### Testing an SCC Application in a Production Environment

13. At some point it will be necessary to test the application in a production environment. To do so, remove any debugging code from the application, then rebuild the application by performing steps 5 through 8 of this procedure. In

<sup>3</sup> The developer is free to choose a different file name.

<sup>4</sup> The developer is free to choose a different file name for the output file. The first argument is the name of the .EXE file created in step 5.

<sup>5</sup> The developer is free to choose a different file name.

<sup>6</sup> The developer is free to choose a different name for the output file. The first argument is the name of the text file mentioned in the previous footnote.

<sup>7</sup> The first argument is the name of the file created in step 8.

step 5 do not specify the options that incorporate debugging information in the .EXE file. In step 6, only one translation need be performed.

14. Load a production-level copy of CP/Q++ (one that lacks the components that support the debugging of applications) into the development coprocessor using TNPrsss.CLU:

**TKNCLU CLU.LOG PL TNPrsss.CLU**

15. Clear any state saved in nonvolatile memory using the procedure described in step 9.
16. Download the file generated in step 8 to the coprocessor using the procedure described in step 10.
17. Wait for the coprocessor to reboot and start the application.

If changes to the application prove necessary, make them and continue with step 13. If additional debugging is required, reload TPrsss.CLU as indicated in step 4 and continue with step 5.

### **Packaging and Releasing an SCC Application**

18. Reset the development coprocessor using TRSrrss.CLU:

**TKNCLU CLU.LOG PL TRSrrss.CLU**

If it again becomes necessary to use the coprocessor for development, begin with step 2 of this procedure.

19. Install the CCA Support Program on the host, install the CCA application in the coprocessor, and configure the coprocessor as a CCA test node following the instructions in chapters 3, 4, and 5 of the *IBM 4758 PCI Cryptographic Coprocessor CCA Support Program*.
20. Generate three RSA keypairs using TKNSGnr's KEYGEN function:

```
TKNSGnr KEYGEN 2 S3KDEVPP.KEY S3KDEVPU.KEY DFT_SKEL.TKN
TKNSGnr KEYGEN 2 DEVSGNPP.KEY DEVSGNPU.KEY DFT_SKEL.TKN
TKNSGnr KEYGEN 2 DEVPKGPP.KEY DEVPKGPU.KEY DFT_SKEL.TKN
```

The first keypair supplies the key to be saved with the developer's application in segment 3. The second and third keypairs are used by TKNSGnr and TKNPkgr, respectively, to generate digital signatures that CLU uses to verify that IBM has authorized its use.

21. Forward each public key generated in step 20 to IBM. Communicate the hash value of each public key (the hash value is also generated by the commands in step 20) to IBM by way of a separate channel to ensure an adversary has not replaced the developer's public key file with another.

After an appropriate contract has been signed, IBM supplies a file (DEV3.TXT) containing a unique owner identifier for the developer and a file (IBMCLT2.TXT) containing the owner identifier associated with segment 2. IBM also creates

- an emergency signature file (ESIGDEV.SIG) incorporating the developer's owner identifier and segment public key and
- a TKNSGnr input file (DEVrrss.TSK) that loads CP/Q++ into an IBM 4758 shipped from the factory and sets the owner identifier associated with segment 3 to the developer's owner identifier.

Finally, IBM supplies certificates for the TKNSGnr and TKNPkgr public files (DEVSGNPU.CRT and DEVPKGPU.CRT, respectively).



IBM generates a hash of each file (DEV3.HSH, IBMCLT2.HSH, ESIGDEV.HSH, DEVrrrss.HSH, DEVSGNPU.HSH, and DEVPKGGPU.HSH) and forwards the files and the hashes (by way of a separate channel) to the developer.

22. Verify that the hashes supplied by IBM in step 21 match the files they cover using TKNSG NR's HASH\_VER function:

```
TKNSG NR HASH_VER DEV3.HSH DEV3.TXT
TKNSG NR HASH_VER IBMCLT2.HSH IBMCLT2.TXT
TKNSG NR HASH_VER ESIGDEV.HSH ESIGDEV.SIG
TKNSG NR HASH_VER DEVrrrss.HSH DEVrrrss.TSK
TKNSG NR HASH_VER DEVSGNPU.HSH DEVSGNPU.CRT
TKNSG NR HASH_VER DEVPKGGPU.HSH DEVPKGGPU.CRT
```

23. Build a version of the application for release (for example, build without debugging information or debug code and change the value of pAgentID->DeveloperID in any calls to sccSignOn and the value of pRequestBlock->AgentID.DeveloperID in any calls to sccRequest to the number supplied in DEV3.TXT).
24. Create an EMBURN3 command that incorporates the application, IBM's segment 2 owner ID, the developer's owner ID, and the developer's unique keys:

```
TKNSG NR EMBURN3 MYAPP.TSK
  part version description
  DEVSGNPU.CRT DEVSGNPP.KEY
  APP.ROD title revision
  S3KDEVPP.KEY ESIGDEV.SIG
  ibmclt2 dev3
  1 1
  a 0 b 0 c 0 d 0 e 0 0
  x 0 0 65535 0 0
  x 0 0 65535 0 0
```

where *part*, *version*, and *description* supply information that is incorporated into the output file, *title* and *revision* supply information that is downloaded to the coprocessor and stored with the application in segment 3, *ibmclt2* is the number supplied in IBMCLT2.TXT, and *dev3* is the number supplied in DEV3.TXT. See Appendix E, "Using Signer and Packager" on page E-1 for details.

25. Combine the task file IBM supplies with the file containing the EMBURN3 command:

```
TKNPKGR DEVPKGGPU.CRT DEVPKGPP.KEY
  2 DEVrrrss.TSK MYAPP.TSK MYAPP.CLU
  part version description
```

where *part*, *version*, and *description* supply information that is incorporated into the output file. See Appendix E, "Using Signer and Packager" on page E-1 for details.

The file generated in this step can be loaded into an empty IBM 4758 or shipped directly to an end user.



---

## Appendix B. Using CLU

The Coprocessor Load Utility (TKNCLU.EXE) interacts with the coprocessor's ROM-based system software to update software in flash. The Coprocessor Load Utility can also obtain information about the coprocessor or reset the coprocessor.

### Syntax

```
TKNCLU logfile {PL | RS | ST} [coprocessornumber] [clufilename]
```

where

- *logfile* is the name of a file to which CLU writes information about the operation and its results.
- The second argument specifies the operation CLU is to perform. Recognized values are as follows:
  - **PL** - Download a file containing software and/or commands to the coprocessor.
  - **RS** - Reset the coprocessor.
  - **ST** - Print information about the coprocessor and the software it contains.
- More than one coprocessor may be installed in a host. *coprocessornumber* identifies the coprocessor with which CLU is to interact. The default is 0.

The number assigned to a particular coprocessor depends on the order in which information about devices in the system is presented to the device driver by the host operating system. At the present time there is no way to tell *a priori* which coprocessor will be assigned a given number.

- *clufilename* is the name of the file containing software and commands to download to the coprocessor. This name appears only if the **PL** operation is specified.



---

## Appendix C. How to Reboot the IBM 4758

An IBM 4758 can be rebooted in any of several ways:

1. Using CLU's RS command, for example:

**TKNCLU CLU.LOG RS**

2. By stopping the device driver and restarting it, for example:

**net stop cryptont  
net start cryptont**

This has the additional benefit of resynchronizing the device driver.

3. The coprocessor reboots at the conclusion of a CLU command or after DRUID downloads an application.
4. If an application on the host calls sccOpenAdapter and the card needs to be rebooted, the device driver will do so.



---

## Appendix D. Building SCC Applications with Microsoft Developer Studio 97

This appendix describes how to configure Microsoft Developer Studio 97\*\* to ensure the proper compiler and linker options are used to build an SCC application. These instructions apply to Microsoft Developer Studio 97\*\* with Microsoft Visual C++ 5.0.

---

### Required Settings for the Host-Side Portion of an SCC Application

Open the Project Settings dialog (Project/Settings...). The required settings under each tab are as follows:

- **C/C++**
  - Precompiled Headers Category -
    - Check “Not using precompiled headers”
  - Add the paths listed below (adjusted as necessary to account for where the Developer’s Toolkit directory tree is located) to “Additional include directories”. The paths must be added in the order shown:
    - ...*.\scctk\cpqenv\nt\msvcasm*
    - ...*.\scctk\include\host*
    - ...*.\scctk\include\common*
- **Link**
  - Object/library modules: Add *cryptont.lib*
  - Add the path listed below (adjusted as necessary to account for where the Developer’s Toolkit directory tree is located) to “Additional library path”.
    - ...*.\scctk\lib\host\nt\msvcasm*

---

### Required Settings for the Coprocessor-Side Portion of an SCC Application

Open the Project Settings dialog (Project/Settings...). The required settings under each tab are as follows:

- **C/C++ -**
  - General Category -
    - Debug info: Select “C7 Compatible”
  - Precompiled Headers Category -
    - Check “Not using precompiled headers”
  - Preprocessor Category -
    - Check “Ignore standard include paths”

Add the paths listed below (adjusted as necessary to account for where the Developer’s Toolkit directory tree is located) to “Additional include directories”. The paths must be added in the order shown:

    - ...*.\scctk\cpqenv\nt\msvcasm*
    - ...*.\scctk\include\sccl*
    - ...*.\scctk\include\common*

- Project Options -
  - Add "/Gs1000000"
- **Link**
  - General Category -
    - Object/library modules: Remove all libraries listed (for example, *kernel32.lib*). Add *clib.lib*, *cpqlib.lib*, *scclib.lib*, and *smlib.lib*
    - Check "Generate debug info" and "Ignore all default libraries"
    - Uncheck "Link incrementally"
  - Customize Category -
    - Uncheck "Link incrementally"
    - Uncheck " Use program database"
  - Debug Category -
    - Check "Debug info" and select "Microsoft format"
  - Input Category -
    - Check "Ignore all default libraries"
    - Object/library modules: Specify *clib.lib*, *cpqlib.lib*, *scclib.lib*, and *smlib.lib*
    - Add the path listed below (adjusted as necessary to account for where the Developer's Toolkit directory tree is located) to "Additional library path".
      - *...lscctk\lib\sccl\nt\msvcasm*
  - Output Category -
    - Entry-point symbol: Specify "startup"
  - Project Options -
    - Add "/fixed:no"



---

## Appendix E. Using Signer and Packager

This appendix describes the use of the signer and packager utilities and explains why the design of the coprocessor makes these utilities necessary.<sup>1</sup>

---

### Coprocessor Memory Segments and Security

The design for the IBM 4758 PCI Cryptographic Coprocessor was motivated by the need to simultaneously satisfy the following requirements:

1. Code must not be loaded into the coprocessor unless IBM or an agent IBM trusts has authorized the operation.
2. Once loaded into the coprocessor, code must not run or accumulate state unless the environment in which it runs is trustworthy.
3. Agents outside the coprocessor that interact with code running on the coprocessor must be able to verify that the code is legitimate and that the coprocessor is authentic and has not been tampered with.
4. Shipment and configuration of coprocessors and maintenance on and upgrades to code inside a coprocessor must not require trusted couriers or security officers.
5. IBM must not need to examine a developer's code or have any knowledge of a developer's private cryptographic keys in order to make it possible for customers to load the developer's code into a coprocessor and run it.<sup>2</sup>

Toward these ends, the design defines four "segments":

- Segment 0 is ROM and contains one portion of "Miniboot". Miniboot is the most privileged software in the coprocessor and among other things implements the protocols described in this section.
- Segment 1 is flash and contains the other portion of "Miniboot". The division of Miniboot into a ROM portion and a Flash portion preserves flexibility (the Flash portion can be changed if necessary) while guaranteeing a basic level of security (implemented in the ROM portion).
- Segment 2 is flash and usually contains the coprocessor operating system.
- Segment 3 is flash and usually contains one or more coprocessor applications.

Segment 0 obviously cannot be changed. Segment 1 can be changed, but should this prove necessary IBM will provide a file that can be downloaded using CLU to effect the change. A developer need not use commands that affect segment 1. The remainder of this chapter therefore deals with changes to segments 2 and 3.

There are seven pieces of information associated with each segment:

1. The identity of the owner of the segment, that is, the party responsible for the software that is to be loaded into the segment. Owner identifiers are two bytes

---

<sup>1</sup> For a thorough overview of the coprocessor's security goals and a description of the security architecture, refer to *Building a High-Performance, Programmable Secure Coprocessor*, Research Report RC21102 published by the IBM T.J. Watson Research Center in February 1998.

<sup>2</sup> Notice in particular that neither the EMBURN3 nor the REMBURN3 command requires IBM to have a copy of the code in segment 3 or the private key corresponding to the public key associated with segment 3.

long.<sup>3</sup> IBM owns segment 1 and issues an owner identifier to any party that is developing code to be loaded into segment 2. An owner of segment 2 issues an owner identifier to any party that is developing code that is to be loaded into segment 3 under the segment 2 owner's authority (that is, while the segment 2 owner owns segment 2).

2. The public key for the owner of the segment.
3. The contents of the segment (that is, the operating system or coprocessor application).
4. Data stored in battery-backed RAM by the code in the segment.
5. The name of the segment (for example, the name of the coprocessor application).
6. The revision level of the contents of the segment (for example, the version number of the coprocessor application).
7. A flag indicating whether or not data stored in BBRAM by the code in the segment is to be cleared if the contents of a more privileged segment change.

Segment 2 and segment 3 can be in one of the following states, depending on how much of the information associated with the segment has been verified:

- UNOWNED - None of the information associated with the segment has been set (that is, it is all unreliable).
- OWNED\_BUT\_UNRELIABLE - The segment has an owner but the rest of the information associated with the segment is unreliable.
- RELIABLE\_BUT\_UNRUNNABLE - All of the information associated with the segment is reliable but the code in the segment should not be allowed to run.
- RUNNABLE - All of the information associated with the segment is reliable and the code in the segment may be allowed to run.

Miniboot enforces the following rules:<sup>4</sup>

- If segment 2's state changes to UNOWNED for any reason, segment 3's state is also changed to UNOWNED.
- If segment 2's state is not RUNNABLE, segment 3's state cannot be RUNNABLE. If segment 2's state changes from RUNNABLE to OWNED\_BUT\_UNRELIABLE or to RELIABLE\_BUT\_UNRUNNABLE, segment 3's state is changed to RELIABLE\_BUT\_UNRUNNABLE. If segment 2's state changes from RUNNABLE to UNOWNED, segment 3's state is also changed to UNOWNED in accordance with the first rule.
- If a segment is not RUNNABLE, the areas of BBRAM controlled by the segment are cleared (that is, any information an application in the segment may have saved in BBRAM is lost).

If the coprocessor's tamper-detection circuitry detects an attempt to compromise the physical security of the coprocessor, all data in BBRAM is cleared and Miniboot changes segment 2's state to UNOWNED. Certain unusual errors affecting segment 1 or segment 2 can also cause segment 2's state to change to UNOWNED, OWNED\_BUT\_UNRELIABLE, or RELIABLE\_BUT\_UNRUNNABLE.

---

<sup>3</sup> An owner identifier of all zeros is reserved and means "no owner". A developer's owner identifier is not necessarily the same as the "Developer Identifier" the developer uses when registering coprocessor applications as described in the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Interface Reference*.

<sup>4</sup> The rules can be expressed in the following manner: 1) a segment can't be owned if its "parent" isn't owned and 2) a segment can't be RUNNABLE if its parent isn't RUNNABLE.

Miniboot will not transfer control to segment 2 after the coprocessor is rebooted unless segment 2's state is RUNNABLE. The code in segment 2 should not transfer control to an application in segment 3 unless segment 3's state is RUNNABLE.<sup>5</sup>

Miniboot changes the state of a segment in response to certain commands Miniboot receives from the host. Figure E-1 shows the state transitions for segment 2 and Figure E-2 on page E-4 shows the state transitions for segment 3. A file that is downloaded to the coprocessor using CLU essentially contains one or more of the pieces of information associated with a segment and one or more Miniboot commands. The Signer utility generates a file containing a single Miniboot command and the corresponding segment information and digitally signs it so CLU can verify the command was produced by an authorized agent. The Packager utility combines signed commands into a single file so that a single download can perform several Miniboot commands. A developer who makes a change to an application during development must use the Signer and the Packager to create a file that contains the revised application and the necessary commands to load it into segment 3<sup>6</sup> and make that segment RUNNABLE. This may entail replacing an existing copy of the application or loading the application into an empty segment. In like manner, prior to shipment of the completed application one or more files must be created to allow the end user to load the application and run it no matter what state segment 3 is in to begin with.

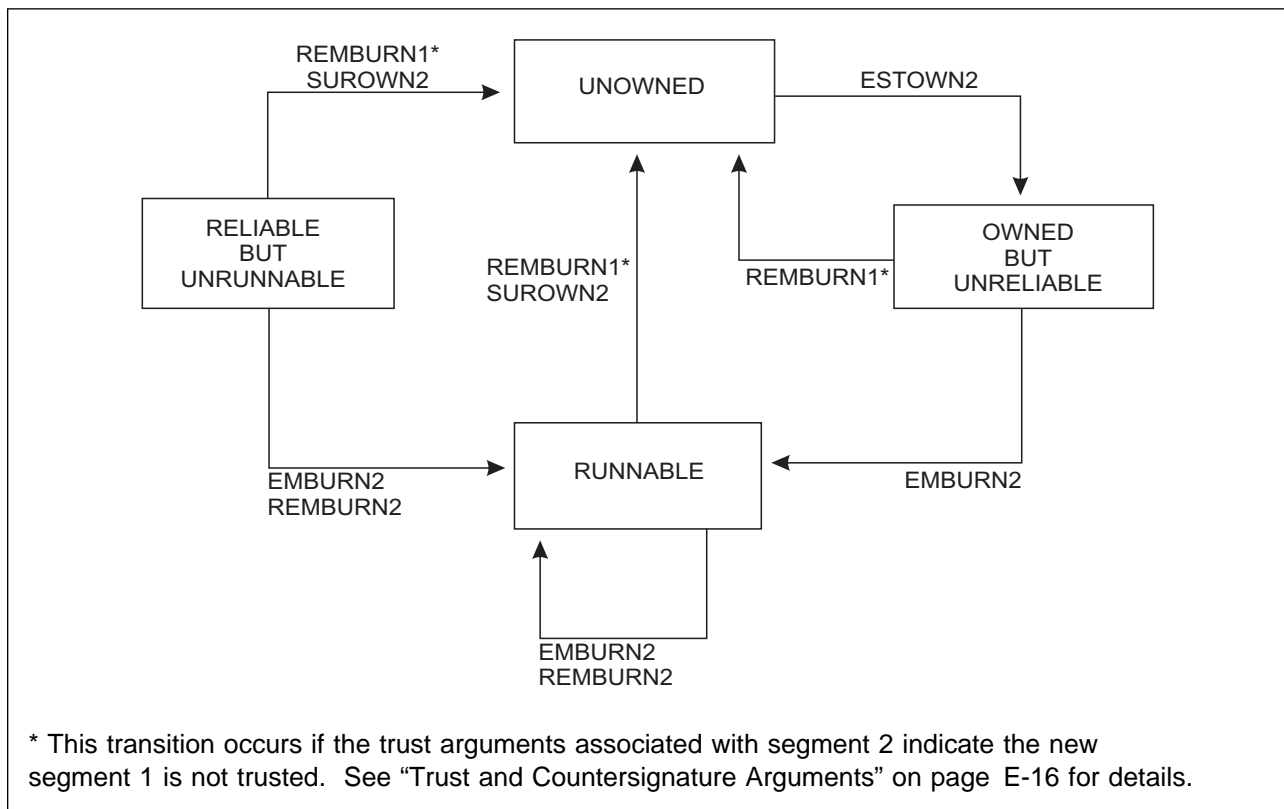


Figure E-1. State Transitions for Segment 2

<sup>5</sup> Segment 3's state is maintained in BBRAM. Information on how to access segment 3's state will appear in the forthcoming Miniboot interface document.

<sup>6</sup> Or segment 2 if the developer is writing an operating system for the coprocessor.

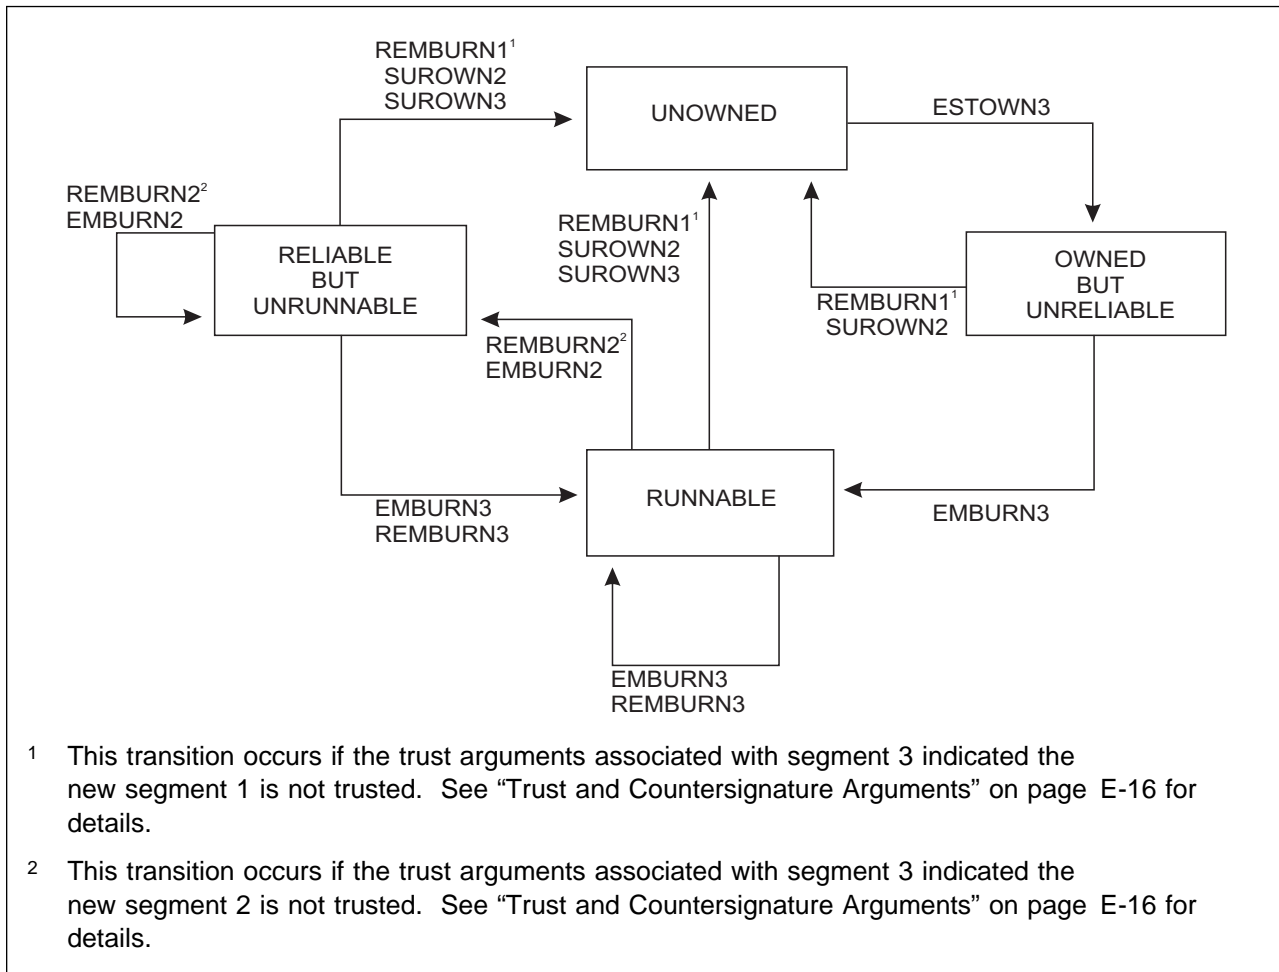


Figure E-2. State Transitions for Segment 3

## The Signer Utility (TKNSGNER.EXE)

The Signer utility (TKNSGNER.EXE) generates a file containing a single Miniboot command and digitally signs it so CLU can verify the command was produced by an authorized agent. The Signer utility also performs certain cryptographic functions. This section describes the syntax of the TKNSGNER command and explains the function of the various TKNSGNER options.

### Syntax

**TKNSGNER** [*function* [*arguments*]] [-F *parm\_file\_name*] [-Q]

The -Q option suppresses all prompts and messages (including error messages). If -Q is specified and TKNSGNER finds it necessary to issue a prompt, the program ends in failure. If the -F option is specified, messages are written to a file named \$SIGNER.RSP.

TKNSGNER reads any arguments that appear on the command line. If no arguments appear on the command line or if the requested function requires more arguments than are specified on the command line, TKNSGNER reads arguments from the file named *parm\_file\_name* if the -F option is specified. Each argument in

the file appears on a separate line. Once the command line and the file specified by the -F option, if present, are exhausted, TKNSG NR issues a prompt for each additional argument required and reads the argument from stdin.

If TKNSG NR reads an argument from stdin, you may select the default for the argument (if there is one) by entering a null line (that is, by pressing the enter key when prompted for the argument), and you must enclose the argument in double quotes if it contains an embedded blank (for example, "This is the description").

## Signer Operations

The first argument to TKNSG NR specifies the Miniboot command TKNSG NR is to generate or the cryptographic function TKNSG NR is to perform and may be one of the following:<sup>7</sup>

### Signer Cryptographic Functions

**KEYGEN**      Generate an RSA key pair.  
**KEYCERT**     Create a certificate for a file containing an RSA public key.  
**HASH\_GEN**    Generate the hash for a file using the SHA1 algorithm.  
**HASH\_VER**    Verify the hash of a file using the SHA1 algorithm.

### Signer Miniboot Command Functions

**EMBURN2**    Load software into segment 2.  
**REMBURN2**   Replace the software in segment 2.  
**SUROWN2**   Surrender ownership of segment 2.  
**ESIG3**       Generate emergency signature for segment 3.  
**ESTOWN3**   Establish ownership of segment 3.  
**EMBURN3**    Load software into segment 3.  
**REMBURN3**   Replace the software in segment 3.  
**SUROWN3**   Surrender ownership of segment 3.

### Signer Miscellaneous Functions

**HELP**        Display instructions about how to use the program.

### Signer IBM-Specific Functions

The following functions are used by IBM to initialize and configure the coprocessor and prepare specific CLU files for developers. Developers writing operating systems or applications for the coprocessor should not need to use these functions (although developers may need to supply as input to the packager files supplied by IBM that direct Miniboot to perform certain of these commands) and they are not otherwise described.

**DATACERT**  
**ESIG2**  
**ESTOWN2**  
**FCVCERT**  
**IBM\_INIT**  
**KEYCERT**  
**RECERT**

<sup>7</sup> Numbers may be used in place of the words listed, as follows: 0 (HELP), 1 (KEYGEN), 2 (HASH\_GEN), 3 (HASH\_VER), 4 (IBM\_INIT), 5 (SIGNFILE), 6 (KEYCERT), 7 (DATACERT), 8 (FCVCERT), 9 (REMBURN1), 10 (REMBURN2), 11 (REMBURN3), 12 (EMBURN2), 13 (EMBURN3), 14 (ESTOWN2), 15 (ESTOWN3), 16 (SUROWN2), 17 (SUROWN3), 18 (ESIG2), 19 (ESIG3), and 20 (RECERT).

## REMBURN1 SIGNFILE

TKNSGMR ignores the case of its first argument (for example, KEYGEN, keygen, and KeyGen are equivalent).

The remainder of this section describes each Signer function, including the arguments it takes, and briefly discusses how it is used during the development process.

## EMBURN2 - Load Software into Segment 2

### Syntax

```
EMBURN2 out_fn filedesc_args sigkey_args image_args privkey_fn esig_fn
        ownid trust1_fl type1_target_args
```

EMBURN2 creates a file that can be downloaded into coprocessor segment 2, which normally contains the coprocessor operating system. The file includes the public key to be associated with segment 2 and the code to load into segment 2. A developer only needs to use this command if the developer is writing an operating system for the coprocessor.

Segment 2 must be owned before an EMBURN2 command can be issued. The file this command causes TKNSGMR to create will often be packaged with commands to ensure the proper agent owns segment 2 (for example, SUROWN2 followed by ESTOWN2). The EMBURN2 command causes the coprocessor to clear data previously stored in BBRAM by code in segment 2 or segment 3.

This command takes the following arguments:

- *out\_fn* is the name of the file TKNSGMR generates to hold the EMBURN2 command. By convention, the file extension is TSK.
- *filedesc\_args* provides certain descriptive information that is incorporated into the output file. See "File Description Arguments" on page E-15 for details.
- *sigkey\_args* specifies the RSA private key that TKNSGMR will use to sign the output file and the certificate provided by IBM for the corresponding RSA public key. See "Signature Key Arguments" on page E-15 for details.
- *image\_args* specifies the name of the file that contains the code to be loaded into segment 2 and provides certain descriptive information about the code that is also downloaded to the coprocessor. See "Image File Arguments" on page E-16 for details.
- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file is the new public key to be associated with segment 2.<sup>8</sup> This key is downloaded to the coprocessor and is used to authenticate subsequent commands that affect segment 2. The key must be the same as the public key contained in the emergency signature information in the *esig\_fn* file.

TKNSGMR includes in the output file a hash of the file enciphered using the private key in the *privkey\_fn* file. The coprocessor uses the public key in the emergency signature information in the *esig\_fn* file to validate the hash and rejects the EMBURN2 command if the validation fails.

---

<sup>8</sup> If desired, the new public key may be the same as the public key currently associated with segment 2, if there is one.

- *esig\_fn* is the name of the file that contains emergency signature information provided by IBM. It includes the public key from the *privkey\_fn* file and includes a hash of the emergency signature information enciphered using the private key corresponding to the public key associated with segment 1. The coprocessor uses the public key associated with segment 1 to validate the hash and rejects the EMBURN2 command if the validation fails.
- *ownid* is the owner identifier currently associated with segment 2. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the EMBURN2 command if the two identifiers are not equal.
- *trust1\_fl* indicates whether or not segment 2s state is to be changed to UNOWNED if the contents of segment 1 change. This flag is downloaded to the coprocessor. See “Trust and Countersignature Arguments” on page E-16 for details.
- *type1\_target\_args* specifies certain conditions that the coprocessor checks before it accepts the new segment 2 information. See “Targeting Arguments” on page E-17 for details.

## EMBURN3 - Load Software into Segment 3

### Syntax

```
EMBURN3 out_fn filedesc_args sigkey_args image_args privkey_fn esig_fn
        seg2_ownid seg3_ownid trust1_fl trust2_fl type2_target_args
```

EMBURN3 creates a file that can be downloaded into coprocessor segment 3, which normally contains a read-only disk image of a coprocessor application. The file includes the public key to be associated with segment 3 and the disk image to load into segment 3.

Segment 3 must be owned before an EMBURN3 command can be issued. The file this command causes TKNSGMR to create will often be packaged with commands to ensure the proper agent owns segment 3 (for example, SUROWN3 followed by ESTOWN3). The EMBURN3 command causes the coprocessor to clear data previously stored in BBRAM by code in segment 3.

This command takes the following arguments:

- *out\_fn* is the name of the file TKNSGMR generates to hold the EMBURN3 command. By convention, the file extension is TSK.
- *filedesc\_args* provides certain descriptive information that is incorporated into the output file. See “File Description Arguments” on page E-15 for details.
- *sigkey\_args* specifies the RSA private key that TKNSGMR will use to sign the output file and the certificate provided by IBM for the corresponding RSA public key. See “Signature Key Arguments” on page E-15 for details.
- *image\_args* specifies the name of the file that is to be loaded into segment 3 (for example, the file that contains the read-only disk image) and provides certain descriptive information about the image that is also downloaded to the coprocessor. See “Image File Arguments” on page E-16 for details.
- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file is the new public key to be associated with segment 3.<sup>9</sup> This key is downloaded to the coprocessor and is used to authenticate subsequent

<sup>9</sup> If desired, the new public key may be the same as the public key currently associated with segment 3, if there is one.

commands that affect segment 3. The key must be the same as the public key contained in the emergency signature information in the *esig\_fn* file.

TKNSGNR includes in the output file a hash of the file enciphered using the private key in the *privkey\_fn* file. The coprocessor uses the public key in the emergency signature information in the *esig\_fn* file to validate the hash and rejects the EMBURN3 command if the validation fails.

- *esig\_fn* is the name of the file that contains emergency signature information provided by IBM. It includes the public key from the *privkey\_fn* file and includes a hash of the emergency signature information enciphered using the private key corresponding to the public key associated with segment 2. The coprocessor uses the public key associated with segment 2 to validate the hash and rejects the EMBURN3 command if the validation fails.
- *seg2\_ownid* is the owner identifier associated with segment 2. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the EMBURN3 command if the two identifiers are not equal.
- *seg3\_ownid* is the owner identifier associated with segment 3. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the EMBURN3 command if the two identifiers are not equal.
- *trust1\_fl* indicates whether or not segment 3's state is to be changed to UNOWNED if the contents of segment 1 change. This flag is downloaded to the coprocessor. See "Trust and Countersignature Arguments" on page E-16 for details.
- *trust2\_fl* indicates whether or not segment 3's state is to be changed to UNOWNED if the contents of segment 2 change. This flag is downloaded to the coprocessor. See "Trust and Countersignature Arguments" on page E-16 for details.
- *type2\_target\_args* specifies certain conditions that the coprocessor checks before it accepts the new segment 3 information. See "Targeting Arguments" on page E-17 for details.

## ESIG3 - Build Emergency Signature for Segment 3

### Syntax

**ESIG3** *out\_fn pubkey\_fn privkey\_fn seg2\_ownid seg3\_ownid type2\_target\_args*

ESIG3 creates a file containing an "emergency signature" that can be provided as an argument to the EMBURN3 command. A developer will only need to use this command if the developer is writing an operating system for the coprocessor: the developer owns segment 2 and uses the ESIG3 command to certify a public key supplied by an agent developing a segment 3 application to run on top of the operating system.

This command takes the following arguments:

- *out\_fn* is the name of the file TKNSGNR generates to hold the emergency signature. By convention, the file extension is BIN.
- *pubkey\_fn* is the name of the file that contains the public key to be associated with segment 3.
- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file must be the public key associated with segment 2. TKNSGNR includes in the output file a hash of the file enciphered using the private key



from the *privkey\_fn* file. The coprocessor uses the public key associated with segment 2 to validate the hash and rejects the EMBURN3 command that contains the emergency signature if the validation fails.

- *seg2\_ownid* is the owner identifier associated with segment 2. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the EMBURN3 command that contains the emergency signature if the two identifiers are not equal.
- *seg3\_ownid* is the owner identifier associated with segment 3. This identifier is assigned by the developer (that is, the segment 2 owner).
- *type2\_target\_args* specifies certain conditions that the coprocessor checks before it accepts the new segment 3 information provided by the EMBURN3 command that contains the emergency signature. See “Targeting Arguments” on page E-17 for details.

## ESTOWN3 - Establish Ownership of Segment 3

### Syntax

```
ESTOWN3 out_fn filedesc_args sigkey_args privkey_fn seg2_ownid seg3_ownid
type2_target_args
```

ESTOWN3 creates a file that directs Miniboot to establish ownership of segment 3, that is, to change segment 3's state from UNOWNED to OWNED\_BUT\_UNRELIABLE. The file includes the owner identifier of the new owner, which is saved in the coprocessor. A developer will only need to use this command if the developer is writing an operating system for the coprocessor: the developer owns segment 2 and uses the ESTOWN3 command to assign ownership of segment 3 to an agent developing a segment 3 application to run on top of the operating system.

Segment 3 must be unowned before an ESTOWN3 command can be issued. The file this command causes TKNSGMR to create will often be packaged with commands to surrender ownership of segment 3 and load software into segment 3 after the new owner is established (for example, SUROWN3 and EMBURN3).

This command takes the following arguments:

- *out\_fn* is the name of the file TKNSGMR generates to hold the ESTOWN3 command. By convention, the file extension is TSK.
- *filedesc\_args* provides certain descriptive information that is incorporated into the output file. See “File Description Arguments” on page E-15 for details.
- *sigkey\_args* specifies the RSA private key that TKNSGMR will use to sign the output file and the certificate provided by IBM for the corresponding RSA public key. See “Signature Key Arguments” on page E-15 for details.
- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file must be the public key associated with segment 2. TKNSGMR includes in the output file a hash of the file enciphered using the private key from the *privkey\_fn* file. The coprocessor uses the public key associated with segment 2 to validate the hash and rejects the ESTOWN3 command if the validation fails.
- *seg2\_ownid* is the owner identifier currently associated with segment 2. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the ESTOWN3 command if the two identifiers are not equal.

- *seg3\_ownid* is the owner identifier to be associated with segment 3. This identifier is assigned by the developer (that is, the segment 2 owner).
- *type2\_target\_args* specifies certain conditions that the coprocessor checks before it accepts the ESTOWN3 command. See “Targeting Arguments” on page E-17 for details.

## HASH\_GEN - Generate Hash for File

### Syntax

**HASH\_GEN** *in\_fn out\_fn*

HASH\_GEN uses the SHA1 algorithm to generate a hash for the file *in\_fn* and writes the result to the file *out\_fn*. The output file consists of groups of five characters representing hexadecimal digits separated by blanks (for example, 03A2 8989 BD90 FFED 0078).

## HASH\_VER - Verify Hash of File

### Syntax

**HASH\_VER** *data\_fn hash\_fn*

HASH\_VER verifies that the hash in the file *hash\_fn* matches the hash the HASH\_GEN function would generate given *data\_fn* as input and issues a message indicating the result (unless the -Q option is specified when TKNSGMR is invoked). The *hash\_fn* file has the same format as the *out\_fn* file generated by the HASH\_GEN function.

## KEYGEN - Generate RSA Key Pair

### Syntax

**KEYGEN** {0 | 2} *keypair\_fn pubkey\_fn skeleton\_fn*

**KEYGEN** 1 *keypair\_fn pubkey\_fn skeleton\_fn transkey\_fn*

**KEYGEN** 3 *pubkey\_fn skeleton\_fn* {0 | 1}

KEYGEN generates an RSA keypair and saves it in the file *keypair\_fn*. The public key is also saved in the file *pubkey\_fn* and the hash of the public key<sup>10</sup> is saved in a file with the same name as *pubkey\_fn* and extension HSH. The file *skeleton\_fn* determines certain characteristics of the keypair, including the key length (that is, the number of bits in the modulus) and the public key exponent. One or more standard skeletons are provided with the Developer’s Toolkit. A developer can also generate customized skeleton files. The file *transkey\_fn* contains an RSA public key.

TKNSGMR uses the PKA\_Key\_Generate CCA verb to generate the keypair. The first argument to KEYGEN determines the *rule\_array* parameter passed with the PKA\_Key\_Generate verb, as follows:

<sup>10</sup> The KEYGEN command computes the hash in the same manner and stores it in the same format as the HASH\_GEN command.

- 0 - Use MASTER for the *rule\_array* parameter. This causes the coprocessor to encrypt the RSA keypair in *keypair\_fn* with the coprocessor CCA master key before returning the keypair.
- 1 - Use XPORT for the *rule\_array* parameter. This causes the coprocessor to encrypt the RSA keypair in *keypair\_fn* with the RSA public key in *transkey\_fn* before returning the keypair.
- 2 - Use CLEAR for the *rule\_array* parameter. This causes the coprocessor to return the RSA keypair in *keypair\_fn* "in the clear" (that is, the file is not encrypted).
- 3 - Use RETAIN for the *rule\_array* parameter. This causes the coprocessor to retain the RSA keypair and not write it to the host. Specify 1 as the last argument if the retained key may be cloned and specify 0 if it may not.

Refer to the *IBM 4758 CCA Basic Services Reference and Guide*, SC31-8609 for details on the format of skeleton files and the PKA\_Key\_Generate CCA verb.

## REMBURN2 - Replace Software in Segment 2

### Syntax

```
REMBURN2 out_fn filedesc_args sigkey_args image_args pubkey_fn privkey_fn
        onnid trust1_fl type2_target_args type3_csign_args
```

REMBURN2 creates a file that can be downloaded into coprocessor segment 2, which normally contains the coprocessor operating system. The file includes the public key to be associated with segment 2 and the code to load into segment 2. A developer will only need to use this command if the developer is writing an operating system for the coprocessor.

Segment 2 must already be occupied (that is, segment 2's state must be RUNNABLE or RUNNABLE\_BUT\_UNRELIABLE) before a REMBURN2 command can be issued.

This command takes the following arguments:

- *out\_fn* is the name of the file TKNSGMR generates to hold the REMBURN2 command. By convention, the file extension is TSK.
- *filedesc\_args* provides certain descriptive information that is incorporated into the output file. See "File Description Arguments" on page E-15 for details.
- *sigkey\_args* specifies the RSA private key that TKNSGMR will use to sign the output file and the certificate provided by IBM for the corresponding RSA public key. See "Signature Key Arguments" on page E-15 for details.
- *image\_args* specifies the name of the file that contains the code to be loaded into segment 2 and provides certain descriptive information about the code that is also downloaded to the coprocessor. See "Image File Arguments" on page E-16 for details.
- *pubkey\_fn* is the name of the file that contains the public key to be associated with segment 2.<sup>11</sup> This key is downloaded to the coprocessor (replacing the key that is already there) and is used to authenticate subsequent commands that affect segment 2.
- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file must be the public key associated with segment 2. TKNSGMR

<sup>11</sup> If desired, the new public key may be the same as the public key currently associated with the segment.

includes in the output file a hash of the file enciphered using the private key from the *privkey\_fn* file. The coprocessor uses the public key associated with segment 2 to validate the hash and rejects the REMBURN2 command if the validation fails.

- *ownid* is the owner identifier associated with segment 2. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the REMBURN2 command if the two identifiers are not equal.
- *trust1\_fl* indicates whether or not segment 2's state is to be changed to UNOWNED if the contents of segment 1 change. See "Trust and Countersignature Arguments" on page E-16 for details.
- *type2\_target\_args* specifies certain conditions that the coprocessor checks before it accepts the new segment 2 information. See "Targeting Arguments" on page E-17 for details.
- *type3\_csign\_args* specifies certain conditions that determine whether or not Miniboot changes segment 3's state to RELIABLE\_BUT\_UNRUNNABLE while updating segment 2.<sup>12</sup> See "Trust and Countersignature Arguments" on page E-16 for details.

## REMBURN3 - Replace Software in Segment 3

### Syntax

```
REMBURN3 out_fn filedesc_args sigkey_args image_args pubkey_fn privkey_fn
seg2_ownid seg3_ownid trust1_fl trust2_fl type3_target_args
```

REMBURN3 creates a file that can be downloaded into coprocessor segment 3, which normally contains a read-only disk image of a coprocessor application. The file includes the public key to be associated with segment 3 and the disk image to load into segment 3.

Segment 3 must already be occupied (that is, segment 3's state must be RUNNABLE or RUNNABLE\_BUT\_UNRELIABLE) before a REMBURN3 command can be issued.

This command takes the following arguments:

- *out\_fn* is the name of the file TKNSGMR generates to hold the REMBURN3 command. By convention, the file extension is TSK.
- *filedesc\_args* provides certain descriptive information that is incorporated into the output file. See "File Description Arguments" on page E-15 for details.
- *sigkey\_args* specifies the RSA private key that TKNSGMR will use to sign the output file and the certificate provided by IBM for the corresponding RSA public key. See "Signature Key Arguments" on page E-15 for details.
- *image\_args* specifies the name of the file that is to be loaded into segment 3 and provides certain descriptive information about the code that is also downloaded to the coprocessor. See "Image File Arguments" on page E-16 for details.
- *pubkey\_fn* is the name of the file that contains the public key to be associated with segment 3.<sup>13</sup> This key is downloaded to the coprocessor (replacing the key that is already there) and is used to authenticate subsequent commands that affect segment 3.

<sup>12</sup> The change to segment 3's state and the updates of segment 2 are performed automatically.

<sup>13</sup> If desired, the new public key may be the same as the public key currently associated with the segment.

- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file must be the public key associated with segment 3. TKNSG NR includes in the output file a hash of the file enciphered using the private key from the *privkey\_fn* file. The coprocessor uses the public key associated with segment 3 to validate the hash and rejects the REMBURN3 command if the validation fails.
- *seg2\_ownd* is the owner identifier associated with segment 2. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the REMBURN3 command if the two identifiers are not equal.
- *seg3\_ownd* is the owner identifier associated with segment 3. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the REMBURN3 command if the two identifiers are not equal.
- *trust1\_fl* indicates whether or not segment 3's state is to be changed to UNOWNED if the contents of segment 1 change. See "Trust and Countersignature Arguments" on page E-16 for details.
- *trust2\_fl* indicates whether or not segment 3's state is to be changed to UNOWNED if the contents of segment 2 change. See "Trust and Countersignature Arguments" on page E-16 for details.
- *type3\_target\_args* specifies certain conditions that the coprocessor checks before it accepts the new segment 3 information. See "Targeting Arguments" on page E-17 for details.

## SUROWN2 - Surrender Ownership of Segment 2

### Syntax

**SUROWN2** *out\_fn filedesc\_args sigkey\_args privkey\_fn ownd type2\_target\_args*

SUROWN2 creates a file that directs Miniboot to surrender ownership of segment 2, that is, to change segment 2's state to UNOWNED.<sup>14</sup> A developer will only need to use this command if the developer is writing an operating system for the coprocessor.

Segment 2 must be owned before a SUROWN2 command can be issued. The file this command causes TKNSG NR to create will often be packaged with commands to grant ownership of segment 2 to another agent and load software into segment 2 (for example, ESTOWN2 followed by EMBURN2).

This command takes the following arguments:

- *out\_fn* is the name of the file TKNSG NR generates to hold the SUROWN2 command. By convention, the file extension is TSK.
- *filedesc\_args* provides certain descriptive information that is incorporated into the output file. See "File Description Arguments" on page E-15 for details.
- *sigkey\_args* specifies the RSA private key that TKNSG NR will use to sign the output file and the certificate provided by IBM for the corresponding RSA public key. See "Signature Key Arguments" on page E-15 for details.
- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file must be the public key associated with segment 2. TKNSG NR includes in the output file a hash of the file enciphered using the private key

<sup>14</sup> This also changes segment 3's state to UNOWNED.

from the *privkey\_fn* file. The coprocessor uses the public key associated with segment 2 to validate the hash and rejects the SUROWN2 command if the validation fails.

- *ownid* is the owner identifier associated with segment 2. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the SUROWN2 command if the two identifiers are not equal.
- *type2\_target\_args* specifies certain conditions that the coprocessor checks before it accepts the SUROWN2 command. See “Targeting Arguments” on page E-17 for details.

## SUROWN3 - Surrender Ownership of Segment 3

### Syntax

```
SUROWN3 out_fn filedesc_args sigkey_args image_args privkey_fn seg2_ownid seg3_ownid
         type3_target_args
```

SUROWN3 creates a file that directs Miniboot to surrender ownership of segment 3, that is, to change segment 3's state to UNOWNED.

Segment 3 must be owned before a SUROWN3 command can be issued. The file this command causes TKNSGMR to create will often be packaged with commands to grant ownership of segment 3 to another agent and load software into segment 3 (for example, ESTOWN3 followed by EMBURN3).

This command takes the following arguments:

- *out\_fn* is the name of the file TKNSGMR generates to hold the SUROWN3 command. By convention, the file extension is TSK.
- *filedesc\_args* provides certain descriptive information that is incorporated into the output file. See “File Description Arguments” on page E-15 for details.
- *sigkey\_args* specifies the RSA private key that TKNSGMR will use to sign the output file and the certificate provided by IBM for the corresponding RSA public key. See “Signature Key Arguments” on page E-15 for details.
- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file must be the public key associated with segment 3. TKNSGMR includes in the output file a hash of the file enciphered using the private key from the *privkey\_fn* file. The coprocessor uses the public key associated with segment 3 to validate the hash and rejects the SUROWN3 command if the validation fails.
- *seg2\_ownid* is the contains the owner identifier. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the REMBURN3 command if the two identifiers are not equal.
- *seg3\_ownid* is the owner identifier associated with segment 3. The coprocessor compares this value to the owner identifier stored in the coprocessor and rejects the REMBURN3 command if the two identifiers are not equal.
- *type3\_target\_args* specifies certain conditions that the coprocessor checks before it accepts the SUROWN3 command. See “Targeting Arguments” on page E-17 for details.

## File Description Arguments

TKNPKGR and many TKNSGMR functions take as arguments certain descriptive information that is incorporated into the files TKNPKGR and TKNSGMR generate. The format of these arguments is as follows:

*partnumber ENumber description*

where

- *partnumber* is a string containing up to eight characters. The string is padded with blanks to the full eight characters before it is incorporated into the output file.
- *ENumber* is a string containing up to eight characters. The string is padded with blanks to the full eight characters before it is incorporated into the output file.
- *description* is a string containing up to 80 characters. The string is padded with blanks to the full 80 characters before it is incorporated into the output file.

*partnumber* is intended to uniquely identify a particular component of a software package (for example, a particular application in a suite). *ENumber* is intended to identify the revision level of the component.

## Signature Key Arguments

TKNSGMR and TKNPKGR incorporate a digital signature in files they generate that are destined to be input to CLU. This allows CLU to verify that the file was generated by an agent authorized to do so by IBM (or by an authority IBM has so authorized).<sup>15</sup> The format of these arguments is

*sigkey\_cert\_fn sigkey\_fn*

where

- *sigkey\_cert\_fn* is the name of the certificate file for the key to be used to sign the output file.
- *sigkey\_fn* is the name of the file containing the RSA private key to be used to sign the output file.

When TKNSGMR creates an output file containing a Miniboot command, TKNSGMR incorporates the certificate from the *sigkey\_cert\_fn* file, computes a hash of the output file, encrypts the hash with the private key in the *sigkey\_fn* file, and appends the encrypted hash to the output file. When CLU processes the file, CLU computes the hash of the relevant portions of the file, extracts the public key from the certificate using the public key corresponding to the private key used to create the certificate<sup>16</sup>, uses the extracted key to decrypt the hash, and verifies that the two hash values match.

<sup>15</sup> The signature key arguments are for the purposes of administrative control. Core security is provided by verification of other signatures and is performed inside the coprocessor.

<sup>16</sup> The public key is compiled into CLU.

## Image File Arguments

Many TKNSG NR functions incorporate an image file (for example, the code that is to be loaded into a segment) into the file TKNSG NR generates. The format of the arguments that apply to an image file is as follows:

*image\_fn title revision*

where

- *image\_fn* is the name of the file to incorporate in the output file.
- *title* is a string containing up to 80 characters. The string is padded with blanks to the full 80 characters before it is incorporated into the output file.
- *revision* is a number between 0 and 65535, inclusive.

*revision* and the last 32 bytes of *title* can be referenced in targeting information. See "Targeting Arguments" on page E-17 for details.

## Trust and Countersignature Arguments

Recall that one of the primary design goals for the IBM 4758 PCI Cryptographic Coprocessor was to ensure that software in the coprocessor must not run or accumulate state unless the environment in which it runs is trustworthy. The use of digital signatures ensures that changes to a segment are authorized (hence trusted) by segments with greater privilege (for example, the initial load of segment 3 must be authorized by the owner of segment 2). But trust operates both ways: changes to a segment that are not trusted by a segment with lesser privilege cause the state of the segment with lesser privilege to become unrunnable (for example, untrusted changes to segment 1 make segment 3 unrunnable).

The TKNSG NR functions that replace the contents of a segment (EMBURN2, EMBURN3, REMBURN2, and REMBURN3) include a flag that indicates how the coprocessor is to change the state of the segment if the contents of a more privileged segment change as a result of a REMBURN command. (Changes caused by an EMBURN command are always untrusted.) See "Coprocessor Memory Segments and Security" on page E-1 for details on segment states. The flag may be 1 (always trust the new more privileged segment), 2 (never trust the new more privileged segment), or 3 (trust the new more privileged segment only if it is countersigned).

If a segment *S* specifies a trust flag of 1 with respect to a more privileged segment *S'*, *S* always trusts changes to *S'*. A REMBURN command that changes the contents of *S'* does not affect the state of *S*.

If a segment *S* specifies a trust flag of 2 with respect to a more privileged segment *S'*, *S* never trusts changes to *S'*. A REMBURN command that changes the contents of *S'* changes the state of *S* to RELIABLE\_BUT\_UNRUNNABLE or to UNOWNED.<sup>17</sup> Note that an EMBURN command that changes the contents of *S'* causes *S*'s state to change in this manner regardless of the value of the trust flag.

If a segment *S* specifies a trust flag of 3 with respect to a more privileged segment *S'*, *S* trusts changes to *S'* only if the new image of *S'* is countersigned with the private key corresponding to the public key associated with *S*. The coprocessor

<sup>17</sup> See Figure E-1 on page E-3 and Figure E-2 on page E-4.



validates the countersignature and changes the state of S to RELIABLE\_BUT\_UNRUNNABLE or to UNOWNED<sup>18</sup> if the countersignature is incorrect.

REMBURN commands that affect segments other than segment 3 (for example, REMBURN2) must therefore include arguments to supply a countersignature. The format of the countersignature arguments is

```
{NoCSig2 | privkey_fn type2_target_args} {NoCSig3 | privkey_fn type3_target_args}
```

where

- **NoCSig2** indicates there is no countersignature provided by segment 2. This option is only applicable to the REMBURN1 command and must be specified exactly as shown (that is, case is important)
- **NoCSig3** indicates there is no countersignature provided by segment 3. This option applies to the REMBURN1 and REMBURN2 commands and must be specified exactly as shown (that is, case is important).
- *privkey\_fn* is the name of a file that contains an RSA keypair. The public key in this file must be the public key associated with the segment that requires the countersignature (for example, the public key for segment 3 if *privkey\_fn* appears instead of **NoCSig3**). If *privkey\_fn* appears, the segment providing the key can also provide a set of targeting arguments for the segment providing the key and each more privileged segment. See "Targeting Arguments" or details.

## Targeting Arguments

The TKNSGMR functions that generate Miniboot commands (EMBURN2, EMBURN3, ESIG3, ESTOWN3, REMBURN2, REMBURN3, SUROWN2, and SUROWN3) incorporate information that specifies certain conditions that must be met before the coprocessor will accept and process the command. Because this information can be used to restrict a command so that it can only be used with coprocessors that already contain certain software or even with a specific individual coprocessor, it is called "targeting information". The format of the arguments that specify targeting information is

```
RTCid RTCid_mask VPDserno VPDserno_mask VPDpartno VPDpartno_mask VPDecno  
VPDecno_mask VPDflags VPDflags_mask bootcount_fl [bootcount_lef[bootcount_right]]  
seg1_info [seg2_info[seg3_info]]
```

where

- *RTCid* and *RTCid\_mask* specify a range of permitted values for the serial number incorporated in the coprocessor chip that implements the real-time clock and the battery-backed RAM.<sup>19</sup> Each of these arguments is a string and may contain as many as eight characters. The arguments should have the same length.

Each character in *RTCid\_mask* must be either ASCII 0 or ASCII 1. TKNSGMR uses *RTCid\_mask* to construct an 8-byte hexadecimal number. Each byte in

<sup>18</sup> See Figure E-1 on page E-3 and Figure E-2 on page E-4.

<sup>19</sup> That is, the value `sccGetConfig` returns in `pInfo->AdapterID`. Refer to the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Interface Reference* for details.

the hexadecimal number is set to 0xFF if the corresponding character in *RTCID\_mask* is ASCII 1 and is set to 0x00 otherwise.

TKNSGNR logically ANDs *RTCID* with the hexadecimal number derived from *RTCID\_mask* and passes the result to the coprocessor. The coprocessor logically ANDs the serial number incorporated in the coprocessor's real-time clock chip with the hexadecimal number derived from *RTCID\_mask* and compares the result to the value generated by TKNSGNR. If they are not equal, the Miniboot command that incorporates the targeting information is rejected.

If a command is intended to apply to all possible coprocessors, specify an arbitrary character for *RTCID* and 0 for *RTCID\_mask*.

- *VPDserno* and *VPDserno\_mask* specify a range of permitted values for the coprocessor's IBM serial number.<sup>20</sup> Each of these arguments is a string and may contain as many as eight characters. The arguments should have the same length.

Each character in *VPDserno\_mask* must be either ASCII 0 or ASCII 1. TKNSGNR uses *VPDserno\_mask* to construct an 8-byte hexadecimal number. Each byte in the hexadecimal number is set to 0xFF if the corresponding character in *VPDserno\_mask* is ASCII 1 and is set to 0x00 otherwise.

TKNSGNR logically ANDs *VPDserno* with the hexadecimal number derived from *VPDserno\_mask* and passes the result to the coprocessor. The coprocessor logically ANDs the coprocessor's IBM serial number with the hexadecimal number derived from *VPDserno\_mask* and compares the result to the value generated by TKNSGNR. If they are not equal, the Miniboot command that incorporates the targeting information is rejected.

If a command is intended to apply to all possible coprocessors, specify an arbitrary character for *VPDserno* and 0 for *VPDserno\_mask*.

- *VPDpartno* and *VPDpartno\_mask* specify a range of permitted values for the coprocessor's IBM part number.<sup>21</sup> Each of these arguments is a string and may contain as many as seven characters. The arguments should have the same length.

Each character in *VPDpartno\_mask* must be either ASCII 0 or ASCII 1. TKNSGNR uses *VPDpartno\_mask* to construct a 7-byte hexadecimal number. Each byte in the hexadecimal number is set to 0xFF if the corresponding character in *VPDpartno\_mask* is ASCII 1 and is set to 0x00 otherwise.

TKNSGNR logically ANDs *VPDpartno* with the hexadecimal number derived from *VPDpartno\_mask* and passes the result to the coprocessor. The coprocessor logically ANDs the coprocessor's IBM part number with the hexadecimal number derived from *VPDpartno\_mask* and compares the result to the value generated by TKNSGNR. If they are not equal, the Miniboot command that incorporates the targeting information is rejected.

If a command is intended to apply to all possible coprocessors, specify an arbitrary character for *VPDpartno* and 0 for *VPDpartno\_mask*.

<sup>20</sup> That is, the value `sccGetConfig` returns in `pInfo->VPD.sn`. Refer to the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Interface Reference* for details.

<sup>21</sup> That is, the value `sccGetConfig` returns in `pInfo->VPD.pn`. Refer to the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Interface Reference* for details.

- *VPDecno* and *VPDecno\_mask* specify a range of permitted values for the coprocessor's IBM engineering change level.<sup>22</sup> Each of these arguments is a string and may contain as many as seven characters. The arguments should have the same length.

Each character in *VPDecno\_mask* must be either ASCII 0 or ASCII 1. TKNSGNR uses *VPDecno\_mask* to construct a 7-byte hexadecimal number. Each byte in the hexadecimal number is set to 0xFF if the corresponding character in *VPDecno\_mask* is ASCII 1 and is set to 0x00 otherwise.

TKNSGNR logically ANDs *VPDecno* with the hexadecimal number derived from *VPDecno\_mask* and passes the result to the coprocessor. The coprocessor logically ANDs the coprocessor's IBM engineering change level with the hexadecimal number derived from *VPDecno\_mask* and compares the result to the value generated by TKNSGNR. If they are not equal, the Miniboot command that incorporates the targeting information is rejected.

If a command is intended to apply to all possible coprocessors, specify an arbitrary character for *VPDecno* and 0 for *VPDecno\_mask*.

- *VPDflags* and *VPDflags\_mask* specify a range of permitted values for the coprocessor's VPD flags.<sup>23</sup> Each of these arguments is a string and may contain as many as 32 characters. The arguments should have the same length.

Each character in *VPDflags\_mask* must be either ASCII 0 or ASCII 1. TKNSGNR uses *VPDflags\_mask* to construct a 32-byte hexadecimal number. Each byte in the hexadecimal number is set to 0xFF if the corresponding character in *VPDflags\_mask* is ASCII 1 and is set to 0x00 otherwise.

TKNSGNR logically ANDs *VPDflags* with the hexadecimal number derived from *VPDflags\_mask* and passes the result to the coprocessor. The coprocessor logically ANDs the last 32 bytes of the coprocessor's Vital Product Data record with the hexadecimal number derived from *VPDflags\_mask* and compares the result to the value generated by TKNSGNR. If they are not equal, the Miniboot command that incorporates the targeting information is rejected.

If a command is intended to apply to all possible coprocessors, specify an arbitrary character for *VPDflags* and 0 for *VPDflags\_mask*.

- *bootcount\_fl*, *bootcount\_left*, and *bootcount\_right* are used as follows: each time the coprocessor boots, it increments one of two counters. The "left count" is a 16-bit number kept in EEPROM that is zero when the coprocessor leaves the factory and is incremented each time the coprocessor boots in a zeroized state (that is, each time the coprocessor is revived after having cleared memory upon detecting an attempt to compromise the coprocessor's security). The "right count" is a 32-bit number that is zero when the coprocessor leaves the factory and is incremented each time the coprocessor is booted in a non-zeroized state.<sup>24</sup> It is set to zero if the coprocessor detects an attempt to

<sup>22</sup> That is, the value `sccGetConfig` returns in `pInfo->VPD.ec`. Refer to the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Interface Reference* for details.

<sup>23</sup> That is, the value `sccGetConfig` returns in the last sixteen bytes of `pInfo->VPD.reserved`. Refer to the *IBM 4758 PCI Cryptographic Coprocessor Custom Software Interface Reference* for details.

<sup>24</sup> Every boot increments either the left count or the right count, so the full 48-bit boot count always increases with each boot. If incrementing either the left count or the right count would cause the counter to overflow, the boot process halts in error.

compromise the coprocessor's security.<sup>25</sup> *bootcount\_fl*, *bootcount\_left*, and *bootcount\_right* specify a range of permitted values for the left and right counts.

*bootcount\_fl* may be 0, 1, or 2. If *bootcount\_fl* is 0, *bootcount\_left* and *bootcount\_right* do not appear and the Miniboot command that incorporates the targeting information is accepted regardless of the left and right counts.

If *bootcount\_fl* is 1, *bootcount\_left* is compared to the left count. The Miniboot command that incorporates the targeting information is rejected if the left count is greater than *bootcount\_left*. *bootcount\_left* must be between 0 and 65535, inclusive, and *bootcount\_right* does not appear in this case.

If *bootcount\_fl* is 2, *bootcount\_left* is compared to the left count and *bootcount\_right* is compared to the right count. The Miniboot command that incorporates the targeting information is rejected if the left count is greater than *bootcount\_left* or if the left count is equal to *bootcount\_left* and the right count is greater than *bootcount\_right*. Use of both counts in this manner can create a Miniboot command that can be downloaded to the coprocessor only once. *bootcount\_left* must be between 0 and 65535, inclusive, and *bootcount\_right* must be between 0 and 4294967295, inclusive, in this case.

If a command is intended to apply to all possible coprocessors, specify 0 for *bootcount\_fl* and omit *bootcount\_left* and *bootcount\_right*.

- *seg1\_info*, *seg2\_info*, and *seg3\_info* specify a range of permitted values for certain of the information associated with segment 1, segment 2, and segment 3, respectively. The format of *seg1\_info*, *seg2\_info*, and *seg3\_info* is

```
segflags segflags_mask revision_min revision_max hash_fl hash
```

where

- *segflags* and *segflags\_mask* specify a range of permitted values for the last 32 bytes of the segment's name or title (as specified in the EMBURN or REMBURN command that loaded the segment into the coprocessor - see "Image File Arguments" on page E-16 for details). By convention, this portion of the name is used to hold information that specifies the version of the code loaded into the segment. Each of these arguments is a string and may contain as many as 32 characters. The arguments should have the same length.

Each character in *segflags\_mask* must be either ASCII 0 or ASCII 1. TKNSGNR uses *segflags\_mask* to construct a 32-byte hexadecimal number. Each byte in the hexadecimal number is set to 0xFF if the corresponding character in *segflags\_mask* is ASCII 1 and is set to 0x00 otherwise.

The coprocessor logically ANDs *segflags* with the 32-byte hexadecimal number derived from *segflags\_mask*. Both quantities are first extended on the right with binary zeros to a length of 80 bytes if necessary. It then logically ANDs the last 32 bytes of the name associated with the segment (as stored in the coprocessor) with the hexadecimal number derived from *segflags\_mask* and compares the two results. If they are not equal, the Miniboot command that incorporates the targeting information is rejected.

<sup>25</sup> The DRUID utility displays the current left and right counts each time it is run.

If a command is intended to apply to all possible coprocessors, specify an arbitrary character for *segflags* and 0 for *segflags\_mask*.

- *revision\_min* and *revision\_max* specify a range of permitted values for the segment's revision level (as specified in the EMBURN or REMBURN command that loaded the segment into the coprocessor - see "Image File Arguments" on page E-16 for details). Each of these arguments is a number between 0 and 65535, inclusive. *revision\_max* must be greater than or equal to *revision\_min*.

The coprocessor compares the revision level associated with the segment (as stored in the coprocessor) with *revision\_min* and *revision\_max*. If the revision level is less than *revision\_min* or greater than *revision\_max*, the Miniboot command that incorporates the targeting information is rejected.

If a command is intended to apply to all possible coprocessors, specify 0 for *revision\_min* and 65535 for *revision\_max*.

- *hash\_fl* and *hash* specify the segment's contents (that is, the code in the segment). *hash\_fl* may be 0 or 1 and *hash* is either 0 or a string containing 20 characters.

If *hash\_fl* is 1, *hash* must be a string containing 20 characters. Each character must be a hexadecimal digit (that is, ASCII 0 through 9, a through f, or A through F) and is interpreted as a 10-byte hexadecimal number (for example, 0F1E2D3C4B5A69788796 is taken to mean 0x0F1E2D3C4B5A69788796). The coprocessor computes the hash value of the contents of the segment using the SHA1 algorithm and compares the hash to the value specified by *hash*. If the two values are not equal, the Miniboot command that incorporates the targeting information is rejected.

If *hash\_fl* is 0, *hash* must also be 0. The Miniboot command is accepted regardless of the contents of the segment.

If a command is intended to apply to all possible coprocessors, specify 0 for *hash\_fl* and 0 for *hash*.

Only *seg1\_info* appears in "type 1" targeting information. The EMBURN2 command incorporates type 1 targeting information.

*seg1\_info* and *seg2\_info* appear in "type 2" targeting information. The EMBURN3, ESIG3, ESTOWN3, REMBURN2, and SUROWN2 commands incorporate type 2 targeting information.

*seg1\_info*, *seg2\_info*, and *seg3\_info* appear in "type 3" targeting information. The REMBURN3 and SUROWN3 commands incorporate type 3 targeting information, and the REMBURN2 command may include type 3 targeting information in its countersignature.

---

## The Packager Utility (TKNPKGR.EXE)

The packager utility (TKNPKGR.EXE) generates a file containing one or more Miniboot commands (each generated by TKNSGMR) and digitally signs it so CLU can verify the command was produced by an authorized agent. This section describes the syntax of the TKNPKGR command and explains the function of the various TKNPKGR options.

Files generated by TKNSGMR are not suitable for use as input to CLU. They must be processed by TKNPKGR, even if the packager is “packaging together” only one Miniboot command.

## Syntax

### TKNPKGR -H

**TKNPKGR -F** *parm\_file\_name*[-Q]

**TKNPKGR** [*sigkey\_args* [*num\_files* [*in\_fn\_list* [*out\_fn* [*outtype* [*filedesc\_args*]]]]]] [-Q]

TKNPKGR ignores the case of its options (for example, **-H** and **-h** are equivalent). Options may be prefixed with a hyphen or a forward slash (for example, **-Q** and **/Q** are equivalent).

The **-Q** option suppresses all prompts and messages (including error messages). If **-Q** is specified and TKNPKGR finds it necessary to issue a prompt, the program ends in failure.

The first form displays instructions about how to use the program. In addition to **-H** and its equivalents, the program accepts **?**, **-?**, and **/?**.

The second form causes TKNPKGR to read arguments from the file named *parm\_file\_name*. Each argument in the file appears on a separate line. Once the file is exhausted, TKNPKGR issues a prompt for each additional argument required and reads the argument from stdin.

The third form causes TKNPKGR to read arguments from the command line. Once the command line is exhausted, TKNPKGR issues a prompt for each additional argument required and reads the argument from stdin.

If TKNPKGR reads an argument from stdin, you may select the default for the argument (if there is one) by entering a null line (that is, by pressing the Enter key when prompted for the argument), and you must enclose the argument in double quotes if it contains an embedded blank (for example, “This is the description”).

TKNPKGR takes the following arguments:

- *sigkey\_args* specifies the RSA private key that TKNPKGR will use to sign the output file and the certificate provided by IBM for the corresponding RSA public key. See “Signature Key Arguments” on page E-15 for details.
- *num\_files* specifies the number of files (each containing a single Miniboot command) TKNPKGR is to combine into a single image. *num\_files* must be greater than zero.
- *in\_fn\_list* is a list containing the name of each file TKNPKGR is to combine into a single image. The files are added to the image in the order in which they appear in the list.
- *out\_fn* is the name of the file TKNPKGR generates to hold the combined input files. By convention, the file extension is CLU. The default is *fn.clu*, where *fn* is the name of the last file in *in\_fn\_list*.
- *outtype* specifies how the output file is intended to be used. Recognized values are as follows:
  - 2 for segment 1
  - 3 for segment 2

- 4 for segment 3
- 5 for the Hardware Lock Monitor
- 6 for the Function Control Vector
- 7 for a key certificate (KEYCERT)
- 8 for a data certificate (DATACERT)
- 9 for any other image
- 10 for reload segment 1 (REMBURN1)
- 11 for reload segment 2 (REMBURN2)
- 12 for reload segment 3 (REMBURN3)
- 13 for reload segment 2 (EMBURN2)
- 14 for reload segment 3 (EMBURN3)
- 15 for establish ownership of segment 2 (ESTOWN2)
- 16 for establish ownership of segment 3 (ESTOWN3)
- 17 for surrender ownership of segment 2 (SUROWN2)
- 18 for surrender ownership of segment 3 (SUROWN3)
- 19 for recertify the coprocessor (RECERT)

Most values of *outtype* are associated with a single TKNSGMR command, which is shown in parenthesis following the description of the value. For example, specify 12 to package a single TKNSGMR file containing a REMBURN3 command. Specify 9 if the output file will contain more than one Miniboot command.

- *filedesc\_args* provides certain descriptive information that is incorporated into the output file. See “File Description Arguments” on page E-15 for details.





---

## Appendix F. Notices

References in this publication to IBM products, programs, or services do not imply that IBM intends to make these available in all countries in which IBM operates. Any reference to an IBM product, program, or service is not intended to state or imply that only IBM's product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any of IBM's intellectual property rights or other legally protectable rights may be used instead of the IBM product, program, or service. Evaluation and verification of operation in conjunction with other products, programs, or services, except those expressly designated by IBM, are the user's responsibility.

IBM may have patents or pending patent applications covering subject matter in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to the IBM Director of Licensing, IBM Corporation, 500 Columbus Avenue, Thornwood, NY, 10594, USA.

---

### Copying and Distributing Softcopy Files

For online versions of this book, we authorize you to:

- Copy, modify, and print the documentation contained on the media, for use within your enterprise, provided you reproduce the copyright notice, all warning statements, and other required statements on each copy or partial copy.
- Transfer the original unaltered copy of the documentation when you transfer the related IBM product (which may be either machines you own, or programs, if the program's license terms permit a transfer). You must, at the same time, destroy all other copies of the documentation.

You are responsible for payment of any taxes, including personal property taxes, resulting from this authorization.

THERE ARE NO WARRANTIES, EXPRESS OR IMPLIED, INCLUDING THE WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow the exclusion of implied warranties, so the above exclusion may not apply to you.

Your failure to comply with the terms above terminates this authorization. Upon termination, you must destroy your machine readable documentation.

---

## Trademarks

The following terms, denoted by an asterisk (\*) in this publication, are trademarks of the IBM Corporation in the United States or other countries or both:

IBM  
VisualAge

The following terms, denoted by a double asterisk (\*\*) in this publication, are the trademarks of other companies:

Acrobat Reader	Adobe Systems, Inc.
Microsoft	Microsoft Corporation
Microsoft Assembler	Microsoft Corporation
Microsoft Visual C++	Microsoft Corporation
Microsoft Developer Studio 97	Microsoft Corporation
Windows	Microsoft Corporation
Windows NT	Microsoft Corporation

---

## List of Abbreviations and Acronyms

<b>API</b>	application program interface	<b>ISO</b>	International Organization for Standardization
<b>ASCII</b>	American National Standard Code for Information Interchange	<b>MD5</b>	message digest 5 (hashing algorithm)
<b>CCA</b>	Common Cryptographic Architecture	<b>PCI</b>	peripheral component interconnect
<b>CLU</b>	Coprocessor Load Utility	<b>PDF</b>	portable document format
<b>CP/Q</b>	Control Program/Q	<b>RSA</b>	Rivest-Shamir-Adleman (algorithm)
<b>FIPS</b>	Federal Information Processing Standard	<b>SCC</b>	secure cryptographic coprocessor
<b>IBM</b>	International Business Machines	<b>TOD</b>	time-of-day (clock)
<b>ICAT</b>	Interactive Code Analysis Tool	<b>UART</b>	universal asynchronous receiver/transmitters
<b>I/O</b>	input/output	<b>VPD</b>	vital product data
<b>IPL</b>	initial program load		



## Glossary

This glossary includes terms and definitions from the *IBM Dictionary of Computing*, New York: McGraw Hill, 1994. This glossary also includes terms and definitions taken from:

- The *American National Standard Dictionary for Information Systems*, ANSI X3.172-1990, copyright 1990 by the American National Standards Institute (ANSI). Copies may be purchased from the American National Standards Institute, 11 West 42 Street, New York, New York 10036. Definitions are identified by the symbol (A) following the definition.
- The *Information Technology Vocabulary*, developed by Subcommittee 1, Joint Technical Committee 1, of the International Organization for Standardization and the International Electrotechnical Commission (ISO/IEC JTC1/SC1). Definitions of published parts of this vocabulary are identified by the symbol (I) following the definition; definitions taken from draft international standards, committee drafts, and working papers being developed by ISO/IEC JTC1/SC1 are identified by the symbol (T) following the definition, indicating that final agreement has not yet been reached among the participating National Bodies of SC1.

### A

**access.** In computer security, a specific type of interaction between a subject and an object that results in the flow of information from one to the other.

**access control.** Ensuring that the resources of a computer system can be accessed only by authorized users and in authorized ways.

**access method.** A technique for moving data between main storage and input/output devices.

**adapter.** Synonym for *expansion card*.

**agent.** (1) An application that runs within the IBM 4758 PCI Cryptographic Coprocessor. (2) Synonym for *secure cryptographic coprocessor application*.

**American National Standard Code for Information Interchange (ASCII).** The standard code, using a coded character set consisting of seven-bit characters (eight bits including parity check), that is used for information interchange among data processing

systems, data communication systems, and associated equipment. The ASCII set consists of control characters and graphic characters. (A)

**American National Standards Institute (ANSI).** An organization consisting of producers, consumers, and general interest groups that establishes the procedures by which accredited organizations create and maintain voluntary industry standards for the United States. (A)

**ANSI.** American National Standards Institute.

**API.** Application program interface.

**application program interface (API).** A functional interface supplied by the operating system, or by a separate program, that allows an application program written in a high-level language to use specific data or functions of the operating system or that separate program.

**ASCII.** American National Standard Code for Information Interchange.

**authentication.** (1) A process used to verify the integrity of transmitted data, especially a message. (T) (2) In computer security, a process used to verify the user of an information system or protected resource.

**authorization.** (1) In computer security, the right granted to a user to communicate with or make use of a computer system. (T) (2) The process of granting a user either complete or restricted access to an object, resource, or function.

**authorize.** To permit or give authority to a user to communicate with or make use of an object, resource, or function.

### B

**battery-backed random access memory (BBRAM).** Random access memory that uses battery power to retain data while the system is powered off. The IBM 4758 PCI Cryptographic Coprocessor uses BBRAM to store persistent data for SCC applications, as well as the coprocessor device key.

**BBRAM.** Battery-backed random access memory.

**bus.** In a processor, a physical facility along which data is transferred.

## C

**call.** The action of bringing a computer program, a routine, or a subroutine into effect, usually by specifying the entry conditions and jumping to an entry point.

(1) (A)

**card.** (1) An electronic circuit board that is plugged into an expansion slot of a system unit. (2) A plug-in circuit assembly. (3) See also *expansion card*.

**CCA.** Common Cryptographic Architecture.

**ciphertext.** (1) Data that has been altered by any cryptographic process. (2) See also *plaintext*.

**cleartext.** (1) Data that has not been altered by any cryptographic process. (2) Synonym for *plaintext*. (3) See also *ciphertext*.

**CLU.** Coprocessor Load Utility.

**Comm\_Mgr.** Communications Manager.

**Common Cryptographic Architecture (CCA).** A comprehensive set of cryptographic services that furnishes a consistent approach to cryptography on major IBM computing platforms. Application programs can access these services through the CCA application program interface.

**Common Cryptographic Architecture (CCA) API.**

The application program interface used to call Common Cryptographic Architecture functions; it is described in the *IBM 4758 CCA Basic Services Reference and Guide*, SC31-8609.

**Communications Manager (Comm\_Mgr).** A CP/Q++ extension for the IBM 4758 PCI Cryptographic Coprocessor that manages communication among the host device driver, SCC applications, and CP/Q++. It handles the receipt and delivery of request headers, and the inbound and outbound data buffers.

**Control Program/Q (CP/Q).** The operating system embedded within the IBM 4758 PCI Cryptographic Coprocessor. The version of CP/Q used by the coprocessor—including extensions to support cryptographic and security-related functions—is known as CP/Q++.

**coprocessor.** (1) A supplementary processor that performs operations in conjunction with another processor. (2) A microprocessor on an expansion card that extends the address range of the processor in the host system, or adds specialized instructions to handle a particular category of operations; for example, an I/O coprocessor, math coprocessor, or a network coprocessor.

**Coprocessor Load Utility (CLU).** A program used to load validated code into the IBM 4758 PCI Cryptographic Coprocessor.

**CP/Q.** Control Program/Q.

**Cryptographic Coprocessor (IBM 4758).** An expansion card that provides a comprehensive set of cryptographic functions to a workstation.

**cryptographic node.** A node that provides cryptographic services such as key generation and digital signature support.

**cryptography.** (1) The transformation of data to conceal its meaning. (2) In computer security, the principles, means, and methods used to so transform data.

## D

**data encrypting key.** (1) A key used to encipher, decipher, or authenticate data. (2) Contrast with *key-encrypting key*.

**Data Encryption Standard Manager (DES\_Mgr).** A CP/Q++ extension that manages the IBM 4758 PCI Cryptographic Coprocessor DES processing hardware.

**decipher.** (1) To convert enciphered data into clear data. (2) Contrast with *encipher*.

**DES\_Mgr.** Data Encryption Standard Manager.

**device driver.** (1) A file that contains the code needed to use an attached device. (2) A program that enables a computer to communicate with a specific peripheral device; for example, a printer, videodisc player, or a CD drive.

## E

**encipher.** (1) To scramble data or convert it to a secret code that masks its meaning. (2) Contrast with *decipher*.

**enciphered data.** (1) Data whose meaning is concealed from unauthorized users or observers. (2) See also *ciphertext*.

**expansion board.** Synonym for *expansion card*.

**expansion card.** A circuit board that a user can plug into an expansion slot to add memory or special features to a computer.

**expansion slot.** One of several receptacles in a PC or RS/6000 machine into which a user can install an expansion card.

## F

**feature.** A part of an IBM product that can be ordered separately from the essential components of the product.

**Federal Information Processing Standard (FIPS).** A standard that is published by the US National Institute of Science and Technology.

**FIPS.** Federal Information Processing Standard

**flash memory.** A specialized version of erasable programmable read-only memory (EPROM) commonly used to store code in small computers.

## H

**hertz (Hz).** A unit of frequency equal to one cycle per second. **Note:** In the United States, line frequency is 60 Hz, a change in voltage polarity 120 times per second; in Europe, line frequency is 50 Hz, a change in voltage polarity 100 times per second.

**host.** As regards to the IBM 4758 PCI Cryptographic Coprocessor, the workstation into which the coprocessor is installed.

## I

**ICAT.** Interactive Code Analysis Tool.

**initial program load (IPL).** (1) The initialization procedure that causes an operating system to commence operation. (2) The process by which a configuration image is loaded into storage. (3) The process of loading system programs and preparing a system to run jobs.

**inline code.** In a program, instructions that are executed sequentially without branching to routines, subroutines, or other programs.

**input/output (I/O).** (1) Pertaining to input, output, or both. (A) (2) Pertaining to a device, process, or channel involved in data input, data output, or both.

**Interactive Code Analysis Tool (ICAT).** A remote debugger used to debug applications running within the IBM 4758 PCI Cryptographic Coprocessor.

**interface.** (1) A boundary shared by two functional units, as defined by functional characteristics, signal characteristics, or other characteristics as appropriate. The concept includes specification of the connection between two devices having different functions. (T) (2) Hardware, software, or both that links systems, programs, and devices.

**International Organization for Standardization (ISO).** An organization of national standards bodies established to promote the development of standards that facilitate the international exchange of goods and services; also, to foster cooperation in intellectual, scientific, technological, and economic activity.

**intrusion latch.** A software-monitored bit that can be triggered by an external switch connected to a jumper on the IBM 4758 PCI Cryptographic Coprocessor. This latch can be used, for example, to detect when the cover of the coprocessor host workstation has been opened. The intrusion latch does not trigger the destruction of data stored within the coprocessor.

**I/O.** Input/output.

**IPL.** Initial program load.

**ISO.** International Organization for Standardization.

## J

**jumper.** A wire that joins two unconnected circuits.

## K

**key.** In computer security, a sequence of symbols used with an algorithm to encipher or decipher data.

## M

**master key.** In computer security, the top-level key in a hierarchy of KEKs.

**miniboot.** Software within the IBM 4758 PCI Cryptographic Coprocessor designed to initialize the CP/Q++ operating system and to control updates to flash memory.

**multi-user environment.** A computer system that supports terminals and keyboards for more than one user at the same time.

## N

**National Institute of Science and Technology (NIST).** Current name for the US National Bureau of Standards.

**NIST.** National Institute of Science and Technology.

**node.** (1) In a network, a point at which one or more functional units connects channels or data circuits. (I) (2) The endpoint of a link or junction common to two or more links in a network. Nodes can be processors, communication controllers, cluster controllers, or

terminals. Nodes can vary in routing and other functional capabilities.

**NT.** See *Windows NT*.

## P

**passphrase.** In computer security, a string of characters known to the computer system and to a user; the user must specify it to gain full or limited access to the system and to the data stored therein.

**private key.** (1) In computer security, a key that is known only to the owner and used with a public key algorithm to decipher data. Data is enciphered using the related public key. (2) Contrast with *public key*. (3) See also *public key algorithm*.

**procedure call.** In programming languages, a language construct for invoking execution of a procedure. (1) A procedure call usually includes an entry name and the applicable parameters.

**public key.** (1) In computer security, a key that is widely known and used with a public key algorithm to encipher data. The enciphered data can be deciphered only with the related private key. (2) Contrast with *private key*. (3) See also *public key algorithm*.

**Public Key Algorithm Manager (PKA\_Mgr).** A CP/Q++ extension that manages the IBM 4758 PCI Cryptographic Coprocessor PKA processing hardware.

## R

**Random Number Generator Manager (RNG\_Mgr).** A CP/Q++ extension that manages the IBM 4758 PCI Cryptographic Coprocessor hardware-based random number generator.

**reduced instruction set computer (RISC).** A computer that processes data quickly by using only a small, simplified instruction set.

**return code.** (1) A code used to influence the execution of succeeding instructions. (A) (2) A value returned to a program to indicate the results of an operation requested by that program.

**RNG\_Mgr.** Random Number Generator Manager.

**RSA algorithm.** A public key encryption algorithm developed by R. Rivest, A. Shamir, and L. Adleman.

## S

**SCC.** Secure cryptographic coprocessor.

**SCC\_Mgr.** Secure Cryptographic Coprocessor Manager.

**secure cryptographic coprocessor (SCC).** An alternate name for the IBM 4758 PCI Cryptographic Coprocessor. The abbreviation "SCC" is used within the product software code.

**secure cryptographic coprocessor (SCC) application.** (1) An application that runs within the IBM 4758 PCI Cryptographic Coprocessor. (2) Synonym for *agent*.

**Secure Cryptographic Coprocessor Manager (SCC\_Mgr).** A CP/Q++ extension that provides high-level management of all agents running within a IBM 4758 PCI Cryptographic Coprocessor. As the "traffic cop", the SCC\_Mgr identifies agents and controls the delivery of their messages and data.

**security.** The protection of data, system operations, and devices from accidental or intentional ruin, damage, or exposure.

**SMIT.** System Management Interface Tool.

**system administrator.** The person at a computer installation who designs, controls, and manages the use of the computer system.

**System Management Interface Tool (SMIT).** An AIX utility program used to maintain the system in good working order and to modify the system to meet changing requirements.

## T

**time-of-day (TOD) clock.** A hardware feature that is incremented once every microsecond, and provides a consistent measure of elapsed time suitable for indicating date and time. The TOD clock runs regardless of whether the processing unit is in a running, wait, or stopped state.

**throughput.** (1) A measure of the amount of work performed by a computer system over a given period of time; for example, number of jobs-per-day. (A) (1) (2) A measure of the amount of information transmitted over a network in a given period of time; for example, a network data-transfer-rate is usually measured in bits-per-second.

**TOD clock.** Time-of-day clock.



## U

**utility program.** A computer program in general support of computer processes. (T)

## V

**verb.** A function possessing an `entry_point_name` and a fixed-length parameter list. The procedure call for a verb uses the syntax standard to programming languages.

**vital product data (VPD).** A structured description of a device or program that is recorded at the manufacturing site.

**VPD.** Vital product data.

## W

**Windows NT.** A Microsoft operating system for personal computers.

**workstation.** A terminal or microcomputer, usually one that is connected to a mainframe or a network, and from which a user can perform applications.

## Numerics

**IBM 4758.** IBM 4758 PCI Cryptographic Coprocessor.



# Index

## A

agent 1-1  
 assembler switches 3-8  
 authentication, software

## B

BLDRODSK Utility 3-10  
 building SCC applications with Microsoft Developer Studio 97 D-1

## C

C runtime library  
   intrinsic functions 3-6, 3-8  
   modified functions  
   supported functions 3-4  
   unsupported functions 3-5  
 CCA support program 2-1, 3-11  
 certification, IBM  
 CLU  
   See Coprocessor Load Utility (CLU)  
 CLU, using B-1  
 code-signing utility 3-11  
 compiler options  
   IBM VisualAge C++ (VACPP) 3-6  
   Microsoft Visual C++ (MSVC++) 3-7  
 Coprocessor Load Utility (CLU)  
   introduction 1-4  
   return codes  
   SCC application load  
   SCC application replace  
   software validation  
   syntax  
 coprocessor memory segments  
 CP/Q  
   debug version 1-5  
   optimized version 1-5  
 CPQXLT Utility 3-10

## D

development environment  
   road map 3-1  
   toolkit components 1-4  
 development process 1-3  
 development process, overview A-1  
 Device Reload Utility (DRUID)  
   description  
   introduction 1-4  
   syntax

Disk Builder Utility  
   description 3-10  
   introduction 1-4  
   syntax 3-11

## E

emergency certificate

## I

IBM VisualAge C++  
   See VisualAge C++  
 installation  
   CCA support program 2-1  
   Coprocessor Load Utility 2-1  
   Disk Builder Utility 2-1  
   include files 2-1  
   Signer Utility 2-1  
   Toolkit 2-1  
   Translator Utility 2-1  
 intrinsic functions 3-6, 3-8

## L

librarian, options 3-10  
 linker switches 3-9  
 loading disk images

## M

makefiles, sample  
   compiler independent  
   Visual C++  
   VisualAge C++  
 MASM assembler 3-8  
 memory segments, coprocessor  
 Microsoft Visual C++  
   See Visual C++  
 MSVC++  
   See Visual C++

## O

operating system, coprocessor  
   See CP/Q  
 options  
   assembler 3-8  
   compiler 3-6  
   librarian 3-10  
   linker 3-9  
 overview of the development process A-1

**P**

packager, using E-1  
 packaging and releasing an SCC application 5-1  
 preparing the development platform 2-5  
 production environment, testing SCC application 4-1

**R**

Read-Only Disk Builder Utility 3-10  
 rebooting the IBM 4758 C-1  
 release components 1-5  
 RSA key pair

**S**

sample makefiles  
   See makefiles, sample  
 SCC application 1-1  
 Signer Utility  
   description 3-11  
   introduction 1-5  
   signing an SCC application  
   syntax  
 signer, using E-1  
 switches  
   assembler 3-8  
   compiler 3-6  
   librarian 3-10  
   linker 3-9  
 syntax  
   Coprocessor Load Utility  
   Disk Builder Utility 3-11  
   Signer Utility  
   Translator Utility 3-10

**T**

testing an SCC application 4-1  
 toolkit components 1-4  
 Translator Utility  
   description 3-10  
   introduction 1-4  
   syntax 3-10

**U**

using CLU B-1  
 using Signer and Packager E-1  
 utilities  
   Coprocessor Load Utility  
   Disk Builder 3-10  
   Signer 3-11  
   Translator 3-10

**V**

VACPP  
   See VisualAge C++  
 validation, software  
 Visual C++  
   intrinsic functions 3-8  
   sample makefile  
   switches 3-7  
 VisualAge C++  
   intrinsic functions 3-6  
   sample makefile  
   switches 3-6