



Netra™ CT Server Software Developer's Guide

For Netra 810 and Netra 410 Servers

Sun Microsystems, Inc.
www.sun.com

Part No. 819-2744-10
February 2007, Revision A

Submit comments about this document at: <http://www.sun.com/hwdocs/feedback>

Copyright 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

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

This document and the product to which it pertains are distributed under licenses restricting their use, copying, distribution, and decompilation. No part of the product or of this document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, JMX, OpenBoot, Java, Netra, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and in other countries.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and in other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

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

Copyright 2007 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054, Etats-Unis. Tous droits réservés.

Sun Microsystems, Inc. a les droits de propriété intellectuels relatants à la technologie qui est décrit dans ce document. En particulier, et sans la limitation, ces droits de propriété intellectuels peuvent inclure un ou plus des brevets américains énumérés à <http://www.sun.com/patents> et un ou les brevets plus supplémentaires ou les applications de brevet en attente dans les Etats-Unis et dans les autres pays.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a.

Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

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

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, JMX, OpenBoot, Java, Netra, et Solaris sont des marques de fabrique ou des marques déposées de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays.

Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciées de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

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

Contents

1. Programming Environment	1
Netra CT Server	2
Hardware Description	2
Alarm Card	2
Host CPU Board	2
Satellite CPU Boards	3
I/O Boards	3
Hot-Swapping Capabilities	3
Software Description	4
Operating System Specifics	6
Managed Object Hierarchy	6
Processor Management Services	6
Multicomputing Network	7
Platform Information Control Library	7
Management Framework	7
SNMP/MIB Support	7
SNMP Interface	7
RMI Interface	8
Developing Applications Using PMS	8

Developing Applications to Interface with MOH or SNMP	8
Developing Applications to Run on Host or Satellite CPU Boards	8
2. Netra CT System Equipment Models	9
Modeling a Netra CT System	10
Managed Objects	10
Viewing the Equipment Model Hierarchies	12
Netra CT 810 System Equipment Models	12
Netra CT 410 System Equipment Models	17
3. Getting Started With Netra CT Element Management Agent API	21
Before You Begin	22
Netra CT Element Management Agent API	22
Netra CT Agent Security	23
Creating Your Application	24
Purpose of the Application	25
Determining the System Configuration Hierarchy	25
Communicating With the Netra CT Agent	26
Finding the Root Object Name	27
Traversing the Containment Hierarchy From a Node	28
Listening for Notifications	28
Registering a Notification Listener With EFDMBean Instance	29
Managing Alarms	29
Registering a NotificationListener With an AlarmNotificationFilter	30
Using the Default AlarmSeverityProfile	31
Creating Your Own AlarmSeverityProfile	32
Assigning a New AlarmSeverityProfile	33

Configuring the Agent to Drive Alarm Card Alarm Outputs	34
▼ To Set Up and Use Alarm Features	35
Clearing Alarms	36
Software Monitoring	37
4. Netra CT Element Management Agent API	41
Interface Overview	42
Summary of JDMK	42
Viewing the Netra CT Management Agent API Online	44
How the API Sections are Organized	44
Netra CT Management Agent Interfaces and Classes	45
5. Simple Network Management Protocol	49
SNMP Overview	50
Management Information Base	50
Object Identifiers	51
Netra CT System SNMP Representation	52
ENTITY-MIB	52
IF-MIB	54
HOST-RESOURCES-MIB	54
Host Resources Running Software Table	54
Host Resources Installed Software Table	55
SUN-SNMP-NETRA-CT-MIB	55
Netra CT Network Element High-Level Objects	56
Physical Path Termination Point Table	56
Equipment Table	57
Equipment Holder Table	58
Plug-in Unit Table	59
Hardware Unit to Running Software Relationship Table	60

Hardware Unit to Installed Software Relationship Table	60
Alarm Severity Identifier Textual Convention	61
Alarm Severity Profile Table	61
Alarm Severity Table	61
Trap Forwarding Table	62
MIB Notification Types	63
SNMP Traps	64
Understanding the MIB Variable Descriptions	65
Changing Midplane FRU-ID	66
Setting High Temperature Alarms	68
▼ To Set the High Temperature Alarm Severity to Major	69

6. Managed Object Hierarchy Software Modules 73

Software Module Design	74
Software Services	74
Software Module MBeans	75
SoftwareMonitorMBean	75
DaemonMBean	75
SoftwareServiceMBean	76
NfsServiceMBean	76
UfsServiceMbean	76
TcpServiceMBean	76
UdpServiceMBean	77
IpServiceMBean	77
EtherIfStatsMBean	77
CgtpServiceMBean	77
RnfsServiceMBean	77

7. Processor Management Services	79
PMS Software Overview	80
PMS Man Pages	83
PMS Examples	84
8. Solaris Operating System APIs	137
Solaris Operating System PICL Framework	138
PICL Frutree Topology	140
Chassis Node Property Updates	141
ChassisType	141
fru Class Properties	141
port Class Node	142
port Class Properties	143
Common Property Updates	145
GeoAddr	145
StatusTime	145
ConditionTime	145
Temperature Sensor Node State	146
PICL Man Page References	146
Dynamic Reconfiguration Interfaces	148
Reconfiguration Coordination Manager	148
Hot-Swap Support	149
Configuration Administration	150
Programming Temperature Sensors Using the PICL API	151
Programming Watchdog Timers Using the PICL API	153
Displaying FRU-ID Data	156
MCNet Support	159

Glossary 161

Index 167

Figures

FIGURE 1-1	Netra CT Server Software	4
FIGURE 2-1	Partial Hardware Resource Hierarchy	10
FIGURE 2-2	Hardware Resource Hierarchy Showing Managed Object Classes	11
FIGURE 2-3	Rear-Access Netra CT 810 System View From Alarm Card	13
FIGURE 2-4	Rear-Access Netra CT 810 System View From Host CPU Board	14
FIGURE 2-5	Rear-Access Netra CT 810 System Host CPU Board Local View	15
FIGURE 2-6	Rear-Access Netra CT 810 System Satellite CPU Board Local View	16
FIGURE 2-7	Netra CT 810 System Satellite CPU Board Local View	16
FIGURE 2-8	Rear-Access Netra CT 410 Diskful System View From Alarm Card	18
FIGURE 2-9	Rear-Access Netra CT 410 Diskful Local View From Host CPU Board	19
FIGURE 2-10	Rear-Access Netra CT 410 System Satellite CPU Board Local View	20
FIGURE 4-1	Key Components of the Java Dynamic Management Kit	43
FIGURE 5-1	Hardware Resource Hierarchy	53
FIGURE 7-1	Netra CT Software Services	80
FIGURE 7-2	PMS Software Services and Interfaces	81
FIGURE 8-1	PICL Daemon (<code>picld</code>) and Plug-Ins	139

Tables

TABLE 1-1	Netra CT Server Software Overview	5
TABLE 2-1	Managed Object Class Definitions	11
TABLE 3-1	Solaris Packages for Netra CT Developer APIs	22
TABLE 3-2	Example of Alarm Output Mapping	34
TABLE 4-1	Netra CT Management Agent Interfaces	45
TABLE 4-2	Netra CT Management Agent Classes	47
TABLE 5-1	Physical Entity Table	53
TABLE 5-2	SUN-SNMP-NETRA-CT-MIB Netra CT NE High-Level Objects	56
TABLE 5-3	SUN-SNMP-NETRA-CT-MIB Physical Path Termination Point Table	56
TABLE 5-4	SUN-SNMP-NETRA-CT-MIB Equipment Table	57
TABLE 5-5	SUN-SNMP-NETRA-CT-MIB Equipment Holder Table	58
TABLE 5-6	SUN-SNMP-NETRA-CT-MIB Plug-In Unit Table	59
TABLE 5-7	SUN-SNMP-NETRA-CT-MIB Hardware Unit to Running Software Relation Table	60
TABLE 5-8	SUN-SNMP-NETRA-CT-MIB Hardware Unit to Installed Software Relationship Table	60
TABLE 5-9	SUN-SNMP-NETRA-CT-MIB Alarm Severity Identifier Textual Conventions	61
TABLE 5-10	SUN-SNMP-NETRA-CT-MIB Alarm Severity Profile Table	61
TABLE 5-11	SUN-SNMP-NETRA-CT-MIB Alarm Severity Table	62
TABLE 5-12	SUN-SNMP-NETRA-CT-MIB Trap Forwarding Table	62
TABLE 5-13	MIB Notification Types	63
TABLE 5-14	SUN-SNMP-NETRA-CT-MIB Traps	64

TABLE 5-15	RFC1213-MIB Traps	65
TABLE 5-16	MIB Variable Syntax	65
TABLE 6-1	Software Services	74
TABLE 7-1	Processor Management Services Man Pages	83
TABLE 8-1	PICL FRUtree Topology Summary	140
TABLE 8-2	PICL FRU State Value Properties	142
TABLE 8-3	PICL FRU Condition Value Properties	142
TABLE 8-4	Port Class State Values	143
TABLE 8-5	Port Condition Values	144
TABLE 8-6	PortType Property Values	144
TABLE 8-7	State Property Values for Temperature Sensor Node	146
TABLE 8-8	PICL Man Pages	146
TABLE 8-9	PICL Temperature Sensor Class Node Properties	151
TABLE 8-10	PICL Threshold Levels and MOH Equivalents	151
TABLE 8-11	Watchdog Plug-in Interfaces for Netra CT 810 and 410 Server Software	155
TABLE 8-12	Properties Under <code>watchdog-controller</code> Node	155
TABLE 8-13	Properties Under <code>watchdog-timer</code> Node	155

Preface

The *Netra CT Server Software Developer's Guide* contains information for developers writing application software for the Netra™ CT 810 and 410 servers. This manual assumes you are a software developer familiar with UNIX® commands and networking applications.

How This Book Is Organized

[Chapter 1](#) contains an overview of the Netra CT software and lists the requirements for developing software applications for the platform.

[Chapter 2](#) displays the system's various equipment models. The diagrams in this chapter demonstrate how the Netra CT software views the hardware components.

[Chapter 3](#) offers a tutorial in writing applications that interface with the Netra CT server software.

[Chapter 4](#) introduces the application programming interfaces for the Netra CT server including the Netra CT element management agent software.

[Chapter 5](#) describes the Netra CT Simple Network Management Protocol (SNMP) management information base (MIB).

[Chapter 6](#) presents the design of the Netra CT software modules and how they relate to each other.

[Chapter 7](#) provides an overview of the Netra CT Processor Management Services (PMS) software.

[Chapter 8](#) defines the Solaris™ Operating System's Platform Information and Control Library (PICL) software and how you can use it to set the watchdog timer.

For obscure or difficult terminology definition, see the Glossary.

Typographic Conventions

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

* The settings on your browser might differ from these settings.

Shell Prompts

Shell	Prompt
C shell	<i>machine-name%</i>
C shell superuser	<i>machine-name#</i>
Bourne shell and Korn shell	\$
Bourne shell and Korn shell superuser	#

Related Documentation

The Netra CT Server documentation is listed in the following table.

Title	Part Number
<i>Netra CT Server Upgrade Guide</i>	819-2745
<i>Netra CT Server Product Overview</i>	819-2742
<i>Netra CT Server Installation Guide</i>	819-2740
<i>Netra CT Server Service Manual</i>	819-2741
<i>Netra CT Server System Administration Guide</i>	819-2743
<i>Netra CT Server Safety and Compliance Manual</i>	819-2746
<i>Netra CT Server Software Developer's Guide</i>	819-2744
<i>Netra CT Server Release Notes</i>	819-2739
<i>Netra CT Server Release Notes for Lucent Technologies</i>	819-2747

You might want to refer to documentation on the following software for additional information: the Solaris Operating System, OpenBoot™ PROM firmware, and the Netra High Availability (HA) Suite.

Documentation, Support, and Training

Sun Function	URL
Documentation	http://www.sun.com/documentation/
Support	http://www.sun.com/support/
Training	http://www.sun.com/training/

Third-Party Web Sites

Sun is not responsible for the availability of third-party web sites mentioned in this document. Sun does not endorse and is not responsible or liable for any content, advertising, products, or other materials that are available on or through such sites or resources. Sun will not be responsible or liable for any actual or alleged damage or loss caused by or in connection with the use of or reliance on any such content, goods, or services that are available on or through such sites or resources.

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can submit your comments by going to:

<http://www.sun.com/hwdocs/feedback>

Please include the title and part number of your document with your feedback:

Netra CT Server Software Developer's Guide, part number 819-2744-10

Programming Environment

This chapter provides an overview of the software environment that forms the basis for developing applications for the Netra CT server:

- [“Netra CT Server” on page 2](#)
- [“Hardware Description” on page 2](#)
- [“Software Description” on page 4](#)

Netra CT Server

The Netra CT server system consists of a host CPU, an alarm card which is the nexus of system management, optionally one or several satellite CPUs, and one or several CompactPCI (cPCI) I/O cards. Different software combinations run on each of these elements as is shown in [FIGURE 1-1](#).

Hardware Description

This section provides brief descriptions of the Netra CT server board components and hot-swapping capabilities. See the *Netra CT Server Product Overview (819-2742)* for more information.

Alarm Card

An alarm card is used in the Netra CT 810 and Netra CT 410 servers to control system functions. The board is plugged into slot 8 for the Netra CT 810, and into slot 1 for the Netra CT 410 server. The alarm card has an embedded operating system, and the boot environment is controlled by boot control firmware. Developers use a command-line interface (CLI) to provide an administrative interface to the system. Drawer-level monitoring and control of the system is accomplished through Managed Object Hierarchy (MOH) and Processor Management Service (PMS) software.

Host CPU Board

The host CPU board is the same for both Netra 810 and Netra CT 410 servers. The board is plugged into slot 1 for the Netra CT 810, and into slot 3 for the Netra CT 410 server. The Solaris Operating System runs on these boards. MOH and PMS provide local and drawer-level monitor and control functions.

Satellite CPU Boards

Several satellite CPU cards can occupy the I/O slots and perform normal CPU functions independently. MOH and PMS provide local monitor and control functions.

I/O Boards

One or more cPCI boards can occupy I/O slots. The I/O boards are controlled by the host CPU and the Solaris OS running on the host CPU board.

Hot-Swapping Capabilities

Boards and other field-replaceable units (FRUs) can be swapped while the system is running, depending on whether or not they conform to the Hot-Swap Specification PICMG 2.1 R 2.0. This ability to hot-swap is a feature that is controllable by software if the board itself is hot-swap compliant. For further information on hot-swap issues, see the *Netra CT Server Product Overview* (819-2742), *Netra CT Server System Administration Guide* (819-2744), and *Netra CT Server Service Manual* (819-2741).

Software Description

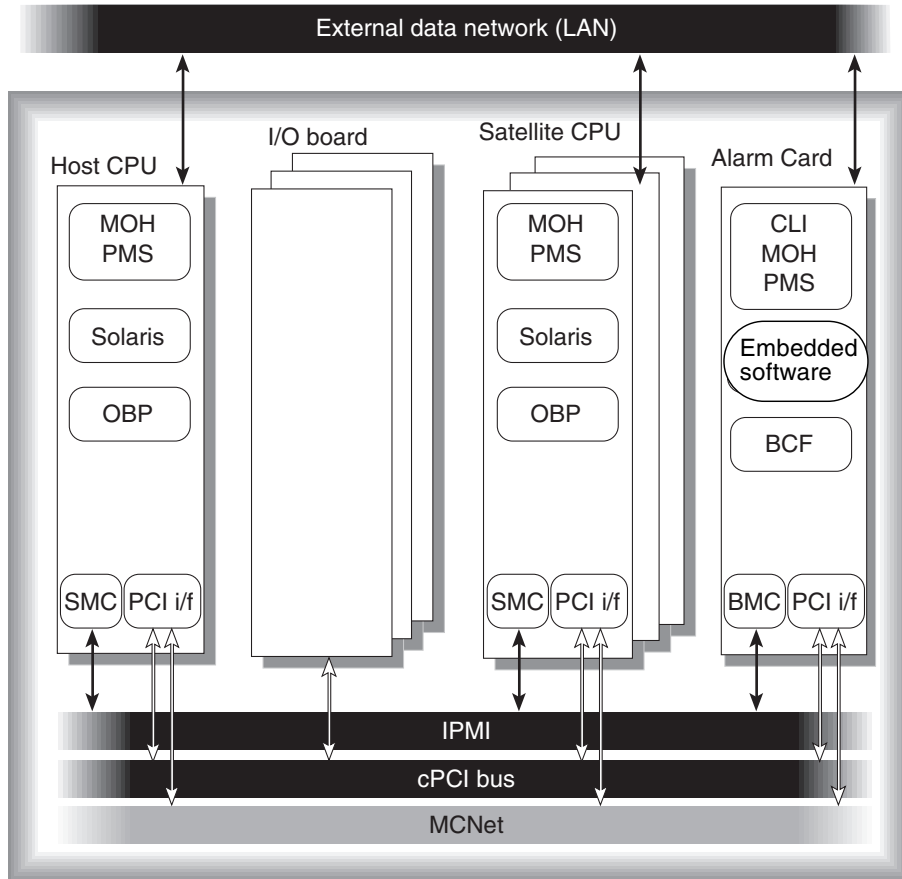


FIGURE 1-1 Netra CT Server Software

The abbreviations shown in [FIGURE 1-1](#) are identified in [TABLE 1-1](#).

TABLE 1-1 Netra CT Server Software Overview

Abbreviation	Name	Description
Solaris	Solaris Operating System	Installed by the user. Runs on the host CPU card and on any satellite CPU cards.
Embedded software	Embedded firmware on the alarm card	Factory-installed on the alarm card. Manages all elements of the Netra CT server that are connected to the midplane.
CLI	Command-line interface	The primary user interface to the alarm card.
MOH	Managed Object Hierarchy	Application that manages the hardware and software components of the system.
PMS	Processor Management Service	Manages processor elements used by client applications.
OpenBoot PROM	OpenBoot PROM firmware and diagnostics	Boot firmware and diagnostics on CPU cards.
BCF	Boot control firmware	Firmware on the alarm card to control booting.
BMC	BMC firmware	Baseboard management controller of the IPMI controller on the alarm card, which provides a command nexus between satellite CPU and RMC client during hot swap unconfiguration operations.
SMC	SMC firmware	System management controller firmware is related to IPMI Controller on CPU cards. SMC APIs provide client access to local resources such as temperature sensors, watchdog subsystems, and local I ² C bus devices; and access to IPMI bus devices.
IPMI	IPMI	Intelligent Platform Management Interface is a communication channel over the cPCI backplane.
MCNet	MCNet	MCNet is a PICMG 2.14 communication protocol over the cPCI backplane. It can be used to communicate between the alarm card, the host CPU card, and any satellite CPU cards which are MCNet capable.

Operating System Specifics

The embedded firmware on the alarm card provides chassis management features that support real-time, multi-threaded applications, and POSIX interfaces to support easy porting of POSIX/UNIX (Solaris OS) applications.

Solaris 9 OS on the host and satellite CPU cards provides APIs such as Platform Information Control Library (PICL), reconfiguration coordination manager (RCM), and `cfgadm` (1M), as explained in [Chapter 8](#). The kernel layer interacts with device drivers to control hardware components of the system such as the CPU cards and the I/O boards. These device drivers bind to the kernel using the device driver interfaces (DDI) and driver kernel interfaces (DKI).

Managed Object Hierarchy

The Managed Object Hierarchy (MOH) is a distributed management application that runs on the alarm card, and host and satellite CPUs. MOH on the alarm card provides drawer-level monitoring of the system. MOH on the CPUs, both host and satellite, provides local views of the board on which it runs, and collaborates to provide the status of its components to the MOH on the alarm card. The various MOHs communicate with one another over MCNet. MOH is discussed further in [Chapter 6](#).

Processor Management Services

Processor management services (PMS) software is an extension to the Netra CT platform services software that addresses the requirements of high-availability application frameworks. PMS software enables client applications to manage the operation of the processor CPU board elements within a single Netra CT system or within a cluster of multiple Netra CT systems.

PMS ensures high availability by monitoring a processor element's fault condition, such as OS hangs, deadlock, and panic. The alarm card provides a server-level view showing the state of each CPU card as a plug-in unit. PMS services are enabled separately on the alarm card and on the host CPU. PMS services are discussed further in [Chapter 7](#).

Multicomputing Network

MCNet uses the cPCI backplane on the Netra CT platform to provide Ethernet-like interface to the CPU cards and the alarm card.

Solaris MCNet driver provides standard Data Link Provider Interface (DLPI) v2 interface to higher level protocols and applications. It appears like any other network interface in Solaris when plumbed.

Platform Information Control Library

This Solaris library provides a method for publishing platform-specific information that clients can access in a way that is not specific to the platform. PICL is discussed further in [Chapter 8](#).

Management Framework

The Java™ Dynamic Management Kit development package provides a framework of managed objects and their associated interfaces. SNMP uses a management information base (MIB), which defines managed objects for the elements within the Netra CT server platform. The managed objects are abstract representations of the resources and services within the system. The following interfaces can be used to manage Netra CT system.

SNMP/MIB Support

The `netract` agent supports the following parts of the MIB:

- System group from MIB II
- Interface group from interface MIB
- Physical entity group from ENTITY-MIB

SNMP Interface

The `netract` agent operates on the alarm card, the system host CPU card, and the satellite CPUs in a distributed manner. They all provide the SNMP interface version 2, and Netra CT-specific instrumentation monitoring.

RMI Interface

The `netract` agent uses JDMK service to support common client-server protocols. These include Remote Method Invocation (RMI) which is the mechanism used to support remote, or distributed access to the managed object hierarchy (MOH).

Developing Applications Using PMS

PMS can run on both the alarm card, and host and satellite CPUs. To develop applications that use PMS on the alarm card or host and satellite CPUs, you need Solaris 9 OS, C compiler version, PMS API, and libraries as described in [Chapter 7](#).

Developing Applications to Interface with MOH or SNMP

To develop applications to interface with MOH or SNMP, you need the Solaris 9 OS, Java Virtual Machine and the Java Dynamic Management Kit and the Netra CT agent library. For more information about Java Dynamic Management Kit refer to *Java Dynamic Management Kit 5.1 Tutorial* (816-7609).

Developing Applications to Run on Host or Satellite CPU Boards

To develop applications to run on host or satellite CPU cards you require Solaris 9 OS to access services such as dynamic reconfiguration (DR) framework, and platform information and control library (PICL) API. Standard Solaris tools such as the `cfgadm(1)` command enable service operations such as configuring and unconfiguring system FRUs.

Netra CT System Equipment Models

This chapter provides illustrations of the Netra CT system equipment models, and contains the following sections:

- [“Modeling a Netra CT System” on page 10](#)
- [“Netra CT 810 System Equipment Models” on page 12](#)
- [“Netra CT 410 System Equipment Models” on page 17](#)

Modeling a Netra CT System

Equipment models show how the Netra CT element management agent software views the Netra CT system hardware. Each equipment model presents a Netra CT system in a containment hierarchy of hardware components, with the midplane at the root of the hierarchy. For example, a cPCI slot might contain an alarm card, which in turn will contain a number of Ethernet and serial ports. These relationships extending from the midplane form a hierarchy of hardware resources. This hierarchy is modeled using relationships between managed objects representing the hardware resources.

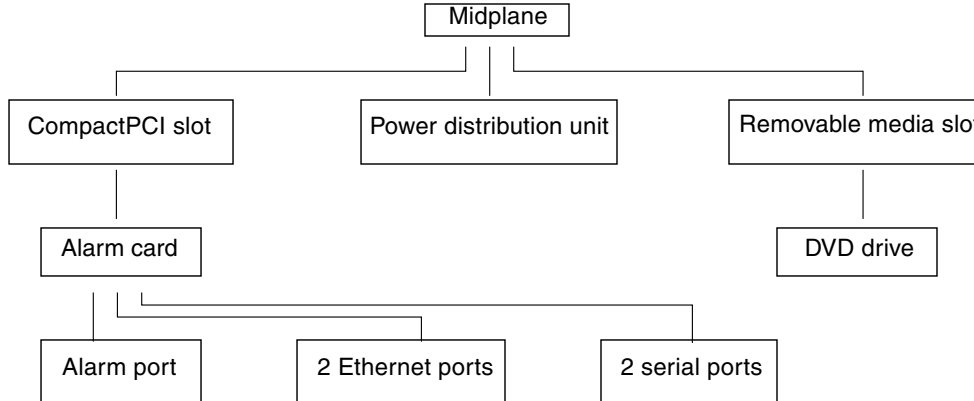


FIGURE 2-1 Partial Hardware Resource Hierarchy

Managed Objects

In the Netra CT software, a managed resource is represented as a managed object, which presents information needed to manage the resource. A managed resource can be represented by a single managed object, or by several managed objects. An agent typically contains or provides views of many managed objects.

[FIGURE 2-2](#) shows the class names of the hardware Netra CT software managed objects, and [TABLE 2-1](#) provides definitions for these objects.

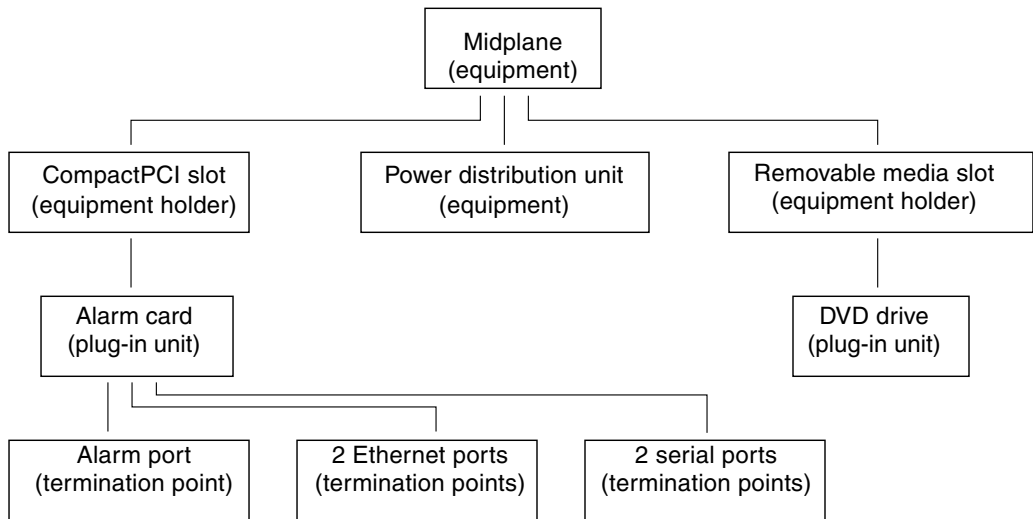


FIGURE 2-2 Hardware Resource Hierarchy Showing Managed Object Classes

TABLE 2-1 Managed Object Class Definitions

Managed Object Class	Definition
Network element	Network elements can be standalone devices or multi-component, geographically distributed systems.
Equipment holder	Represents physical resources of the network element that are capable of holding other physical resources, for example, CompactPCI slots, fan tray slots, and system controller board slots.
Plug-in unit	This managed entity represents equipment that can be physically inserted or removed from slots of the system (for example, CompactPCI I/O cards and power supply units).
Equipment	Equipment represents those externally manageable physical components which are not FRUs (for example, a power distribution unit or a CPU temperature sensor) of a network that are not modeled as a plug-in unit or an equipment holder.
Termination point	Represents the points where physical paths terminate (for example, Ethernet and serial ports) and physical path functions.

Viewing the Equipment Model Hierarchies

Both the SNMP interface and the Java Management Extensions (JMX) compatible Netra CT element management API provide ways to traverse the equipment containment hierarchy. You can view the managed objects of a Netra CT system through the system's alarm card or through the host CPU board. You can also view the managed objects from the agent on any satellite CPU board. In both system-wide views, the system's midplane is at the top of the equipment hierarchy and all other hardware objects (slots, fan trays, I/O cards, and so on) are displayed subordinate to the midplane.

When viewing the system through the alarm card (defined as the *system view from the alarm card*), the alarm card's termination points (alarm port, Ethernet ports, and serial ports) are displayed in the model, but the host CPU board's termination points are not displayed.

Conversely, when you view the system through the host CPU board (the *system view from host CPU board*), the alarm card's termination points are not displayed, but the host CPU board's termination points, and any hardware connected to the host CPU board (for example, SCSI devices), is displayed.

You can also view the equipment model with the host CPU board or a supported satellite CPU board as the network element at the top of the hierarchy. In these models (defined as the *host CPU board local view* and *satellite CPU board local view*), only the objects directly controlled by the host or satellite CPU board are displayed. Other objects, like the midplane, alarm card, and the power distribution unit, are not seen in these equipment models.

[“Netra CT 810 System Equipment Models” on page 12](#) and [“Netra CT 410 System Equipment Models” on page 17](#) present the equipment models for the Netra CT rear-access systems. These sections contain the equipment models shown in the system alarm card view, the host CPU board view, and the host and satellite CPU board views.

Netra CT 810 System Equipment Models

This section discusses the following equipment models of the Netra CT 810 server:

- [“Rear-Access Netra CT 810 System View From Alarm Card” on page 13](#)
- [“Rear-Access Netra CT 810 System View From Host CPU Board” on page 14](#)
- [“Rear-Access Netra CT 810 System Satellite CPU Board Local View” on page 16](#)
- [“Netra CT 810 System Satellite CPU Board Local View” on page 16](#)

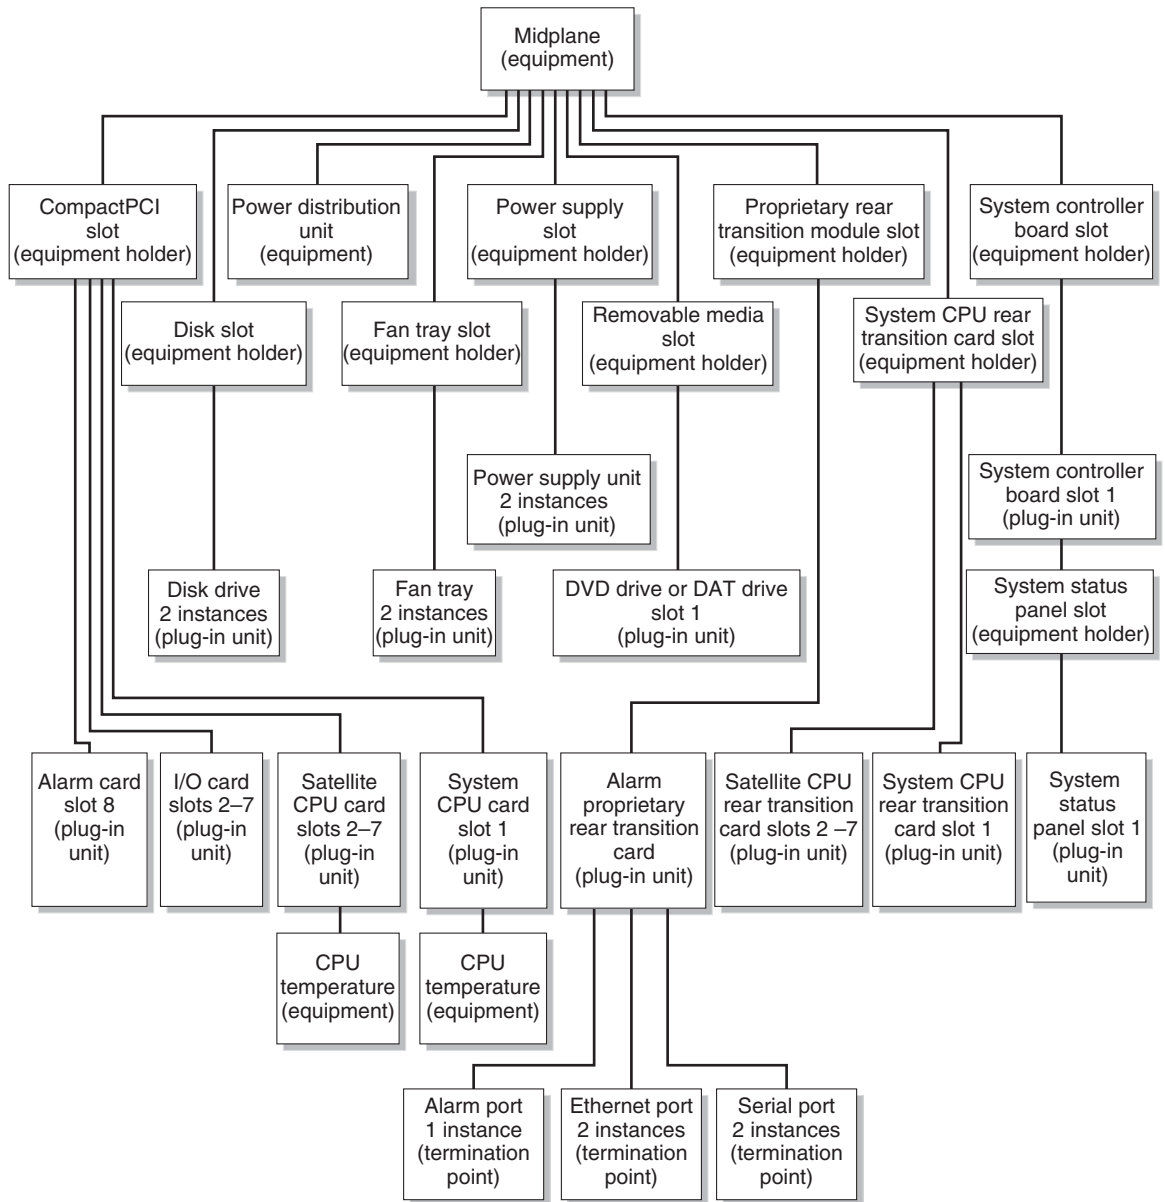


FIGURE 2-3 Rear-Access Netra CT 810 System View From Alarm Card

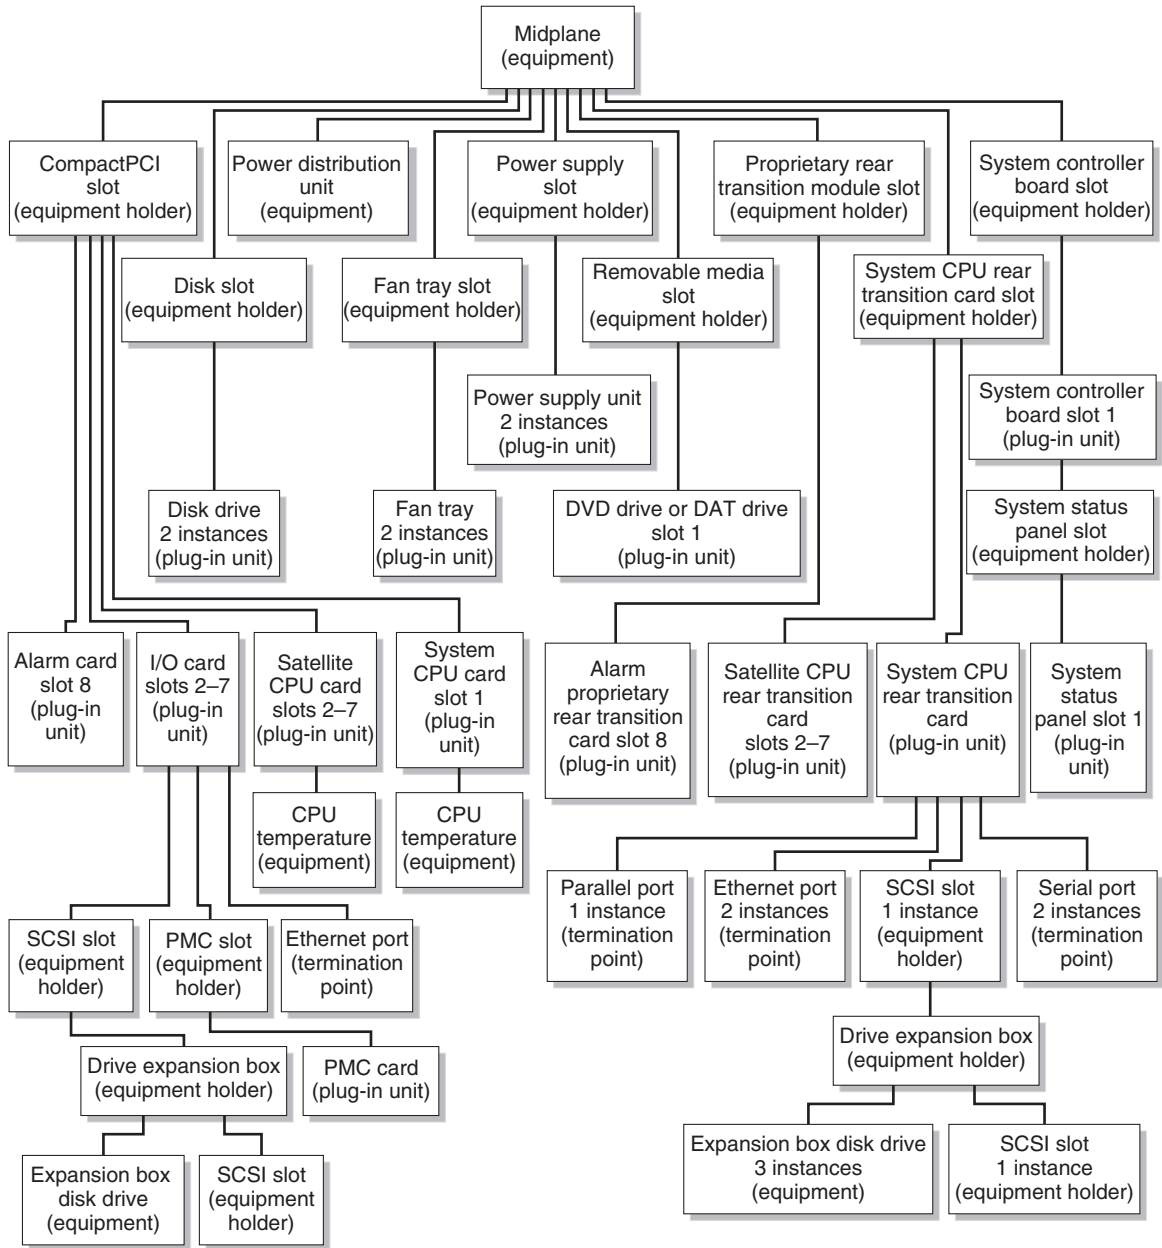


FIGURE 2-4 Rear-Access Netra CT 810 System View From Host CPU Board

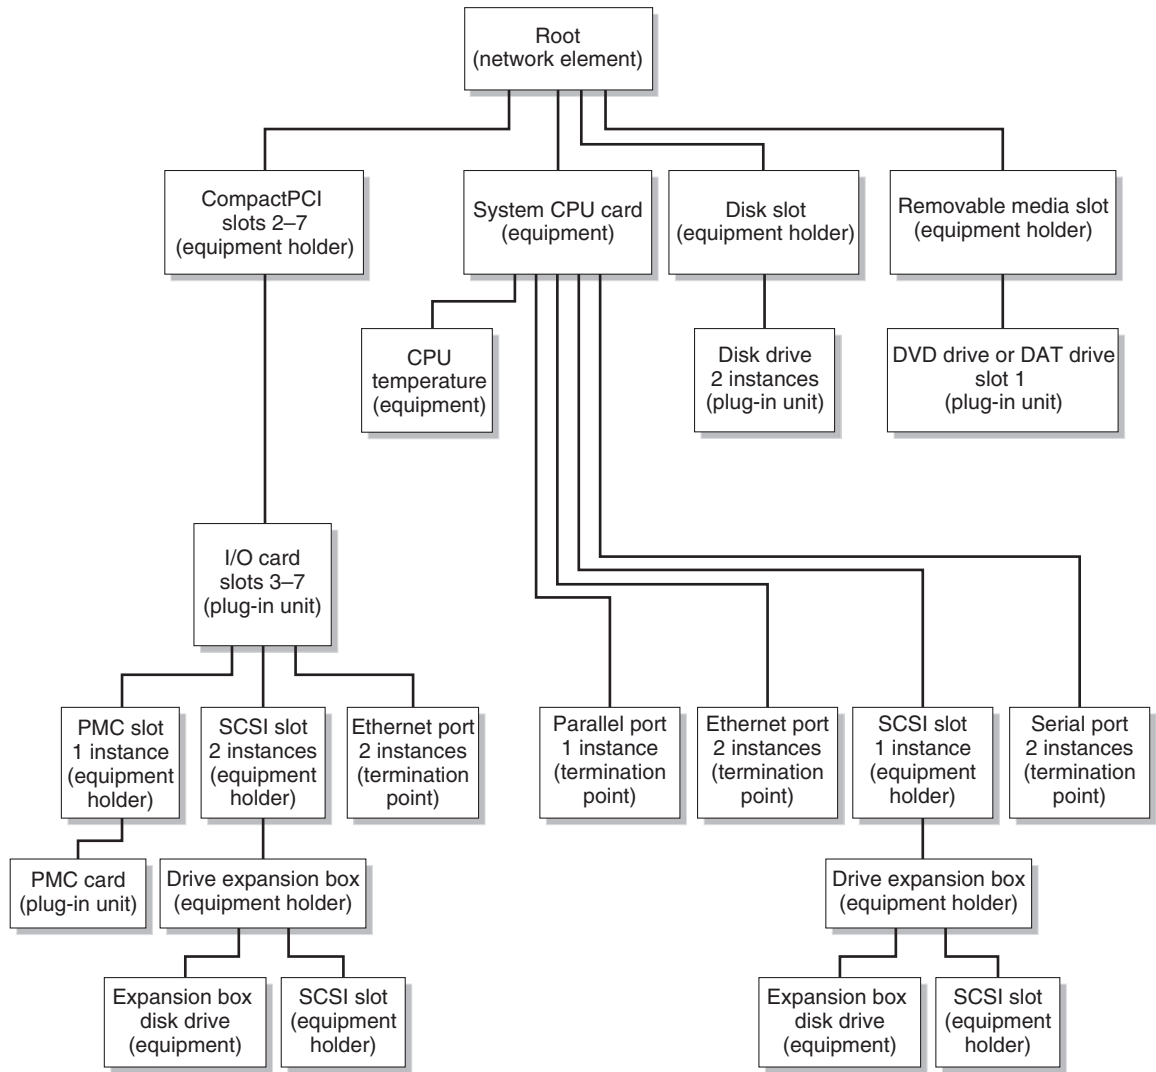


FIGURE 2-5 Rear-Access Netra CT 810 System Host CPU Board Local View

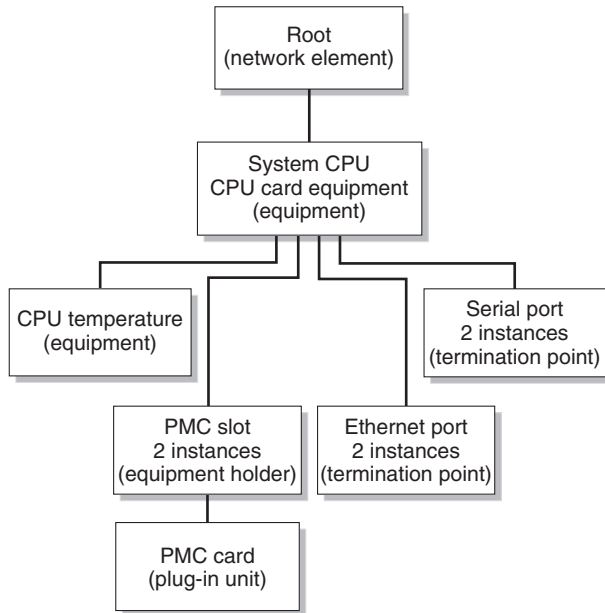


FIGURE 2-6 Rear-Access Netra CT 810 System Satellite CPU Board Local View

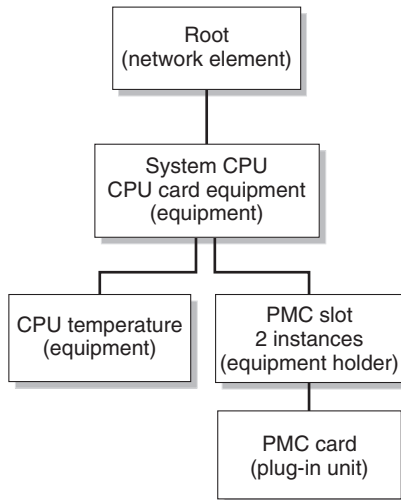


FIGURE 2-7 Netra CT 810 System Satellite CPU Board Local View

Netra CT 410 System Equipment Models

This section discusses the following equipment models for the Netra CT 410 server:

- [“Rear-Access Netra CT 410 Diskful System View From Alarm Card”](#) on page 18
- [“Rear-Access Netra CT 410 Diskful Local View From Host CPU Board”](#) on page 19
- [“Rear-Access Netra CT 410 System Satellite CPU Board Local View”](#) on page 20

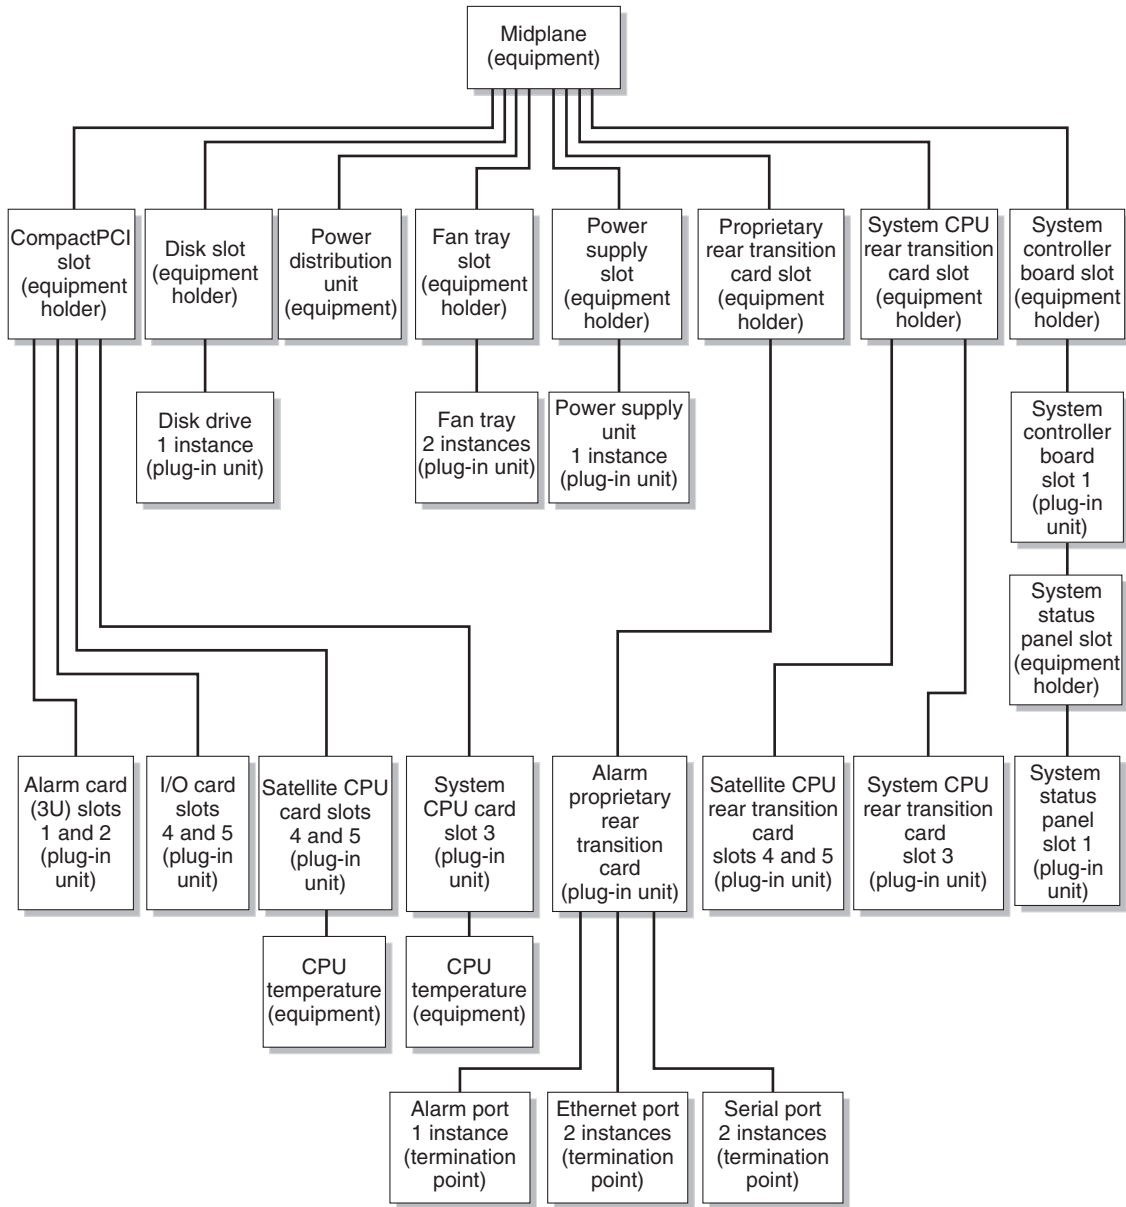


FIGURE 2-8 Rear-Access Netra CT 410 Diskful System View From Alarm Card

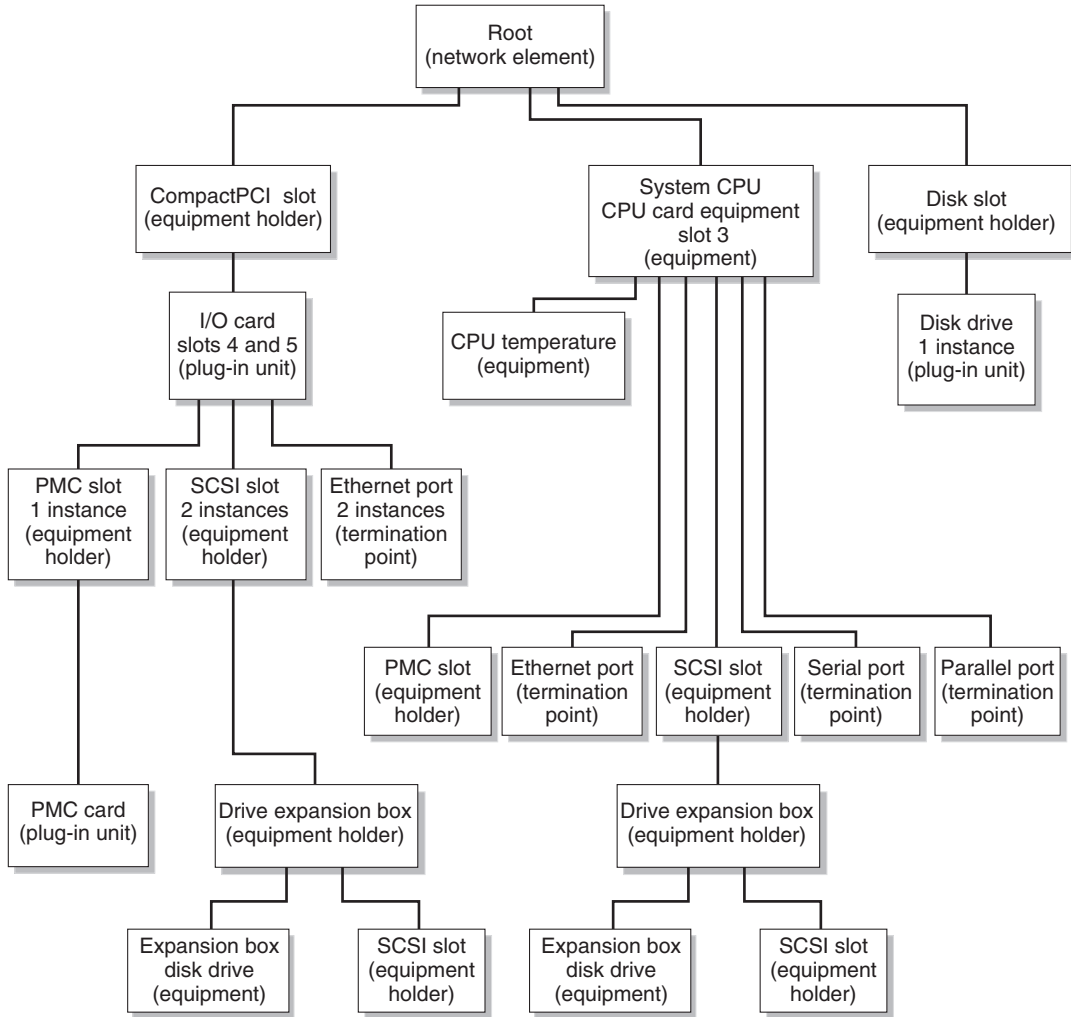


FIGURE 2-9 Rear-Access Netra CT 410 Diskful Local View From Host CPU Board




FIGURE 2-10 Rear-Access Netra CT 410 System Satellite CPU Board Local View

Getting Started With Netra CT Element Management Agent API

This chapter explains how to get started writing applications that interface with the Netra CT element management agent, using the Java Management Extensions (JMX) compatible Java API supported by the Netra CT management agent. The chapter consists of:

- [“Before You Begin” on page 22](#)
- [“Netra CT Element Management Agent API” on page 22](#)
- [“Creating Your Application” on page 24](#)

Before You Begin

You should become acquainted with the topology of the Netra CT server (see Chapter 2), and have some knowledge of Java programming, JMX specifications, and JDMK framework. For more information about JDMK refer to *Java Dynamic Management Kit 4.2 Tutorial* (806-6633), or go to <http://java.sun.com/docs/books/tutorial/index.html>.

Verify that you have the Solaris OS on your development system. In addition, you can download the required Netra CT patch packages from:

<http://sunsolve.com>

These packages consist of:

TABLE 3-1 Solaris Packages for Netra CT Developer APIs

Package	Description
SUNW2jdrct	Java Runtime Java Dynamic Management Kit (JDMK) package
SUNWctmgx	Netra CT management agent package
SUNWctac	Alarm card firmware package that includes the Netra CT management agent

You will use these installed packages to work with this tutorial.

Netra CT Element Management Agent API

The Netra CT server software package includes various modules and extensions (see “[Operating System Specifics](#)” on page 6), and the `netract` agent is one of these.

The `netract` agent, when appropriately invoked, provides configuration monitoring and fault monitoring. This enables you to investigate the installed system, and to determine whether the components are running smoothly.

Individual `netract` agents run on the alarm card, the host CPU board, and any satellite CPU board. A management application must be able to talk to the different agents and gather information about the system into a database.

Each `netract` agent notifies the management application of any changes, such as hardware or software configuration changes, and also detects faults when they occur.

The `netract` agent provides two different interfaces for the management applications, one is SNMP version 2C interface, the other is a JMX compatible Java API called Netra CT management agent API. This chapter provides an introduction on how to write a management application using this Java API.

Netra CT Agent Security

For JMX, JDMK, and RMI connectivity, the Netra CT agent provides security by authenticating the application connecting to it through the context of a valid username and password pair.

The username and password must be previously created in the alarm card database through the alarm card CLI. An account on the alarm card consists of username, password, and permission. The Netra CT agent has only two permissions: read-only and read-write. User account on the alarm card must have ALL PRIVILEGES ENABLED to have read-write permission. (See the *Netra CT Server System Administration Guide* for details on setting up user accounts.)

There is a *security flag* used to enable and disable the security feature. This flag is stored persistently and its default value is false. The security flag can be set to true or false via alarm card CLI command `setmohsecurity`. A reset of alarm card is required after changing the flag for the feature to take effect. (See the *Netra CT Server System Administration Guide* for information on the `setmohsecurity` and `showmohsecurity` CLI commands).

You can get the state of security flag with the alarm card CLI command `showmohsecurity` or with the API by using the NEMBean's `getSecurityFlag` method.

If the flag is `true`, security is on. This means the application that connects to Netra CT agent must provide a valid username and password to be able to establish connection.

If the flag is `false`, security is off and no authentication is done. It does not matter whether an application provides username and password or not, the application is always allowed to connect.

Sample code with Netra CT Security is shown [CODE EXAMPLE 3-1](#).

CODE EXAMPLE 3-1 Sample Code With Netra CT Security

```
...
// set up the authentication info
AuthClient.setAuthInfo(connectorClient, username, password);
// now connect to the agent...
connectorClient.connect();
...
```

Creating Your Application

Creating an application to interface with and manage the configuration of the Netra CT server involves a series of steps. You must be able to:

- Enquire into the hierarchy of the system configuration
- Monitor notifications
- Monitor alarms

1. Cut and paste the relevant code example into a text editor, make any necessary adjustments, and compile the code.

Make sure that SUNW2jdk is installed before trying to compile `Client.java`. Refer to the *Java Dynamic Management Kit 4.2 Tutorial* for background information on `Client.java`.

2. To compile `Client.java`, issue the command `/usr/j2se/bin/javac -classpath:`

```
$ /usr/j2se/bin/javac -classpath \
/opt/SUNWjdkm4/jdkm4.2/1.2/lib/jdkmrt.jar: \
/opt/SUNWnetract/mgmt2.0/lib/agent.jar Client.java
```

Compiling `Client.java` should produce the file `Client.class`. If you have difficulty, refer to the Java Tutorial example of running a simple client.

3. Before running `Client.java`, start the agent by issuing the command `/opt/SUNWnetract/mgmt2.0/bin/ctmgx start`

```
$ /opt/SUNWnetract/mgmt2.0/bin/ctmgx start
```


4. Use the following command to run `Client.java`:

```
$ /usr/j2se/bin/java -classpath \  
.: /opt/SUNWjdmk/jdmk4.2/1.2/lib/jdmkrt.jar: \  
/opt/SUNWnetract/mgmt2.0/lib/agent.jar Client
```

The following sections point out various features of the Netra CT element management API.

Purpose of the Application

First, a management application needs to know how the system is configured. The simplest example sets up an agent describing the hardware containment hierarchy. From the root of this tree, you can develop the management tree to show, for example, how many fans there are, which cards are in which slots and so on. Developing code that begins this action is the purpose of [“Determining the System Configuration Hierarchy” on page 25](#).

[“Listening for Notifications” on page 28](#) deals with developing a way of monitoring notifications such as power on and power off to a particular slot or device.

[“Managing Alarms” on page 29](#) covers alarm management — how to handle the receiving and transmitting of system alarms such as CPU overtemperature alarm.

[“Software Monitoring” on page 37](#) shows how to get a list of running software services and daemons so they can be registered to receive notice of events.

Determining the System Configuration Hierarchy

In this section you develop a client to print out the object names of the MBeans representing the system. A complete description of Mbeans, together with examples, can be found in the JDMK documentation.

1. Ensure that you have the appropriate software installed on the development system for the application you intend to develop.

Refer to the *Netra CT Server System Administration Guide* if you need help in installing the appropriate software.

2. Go to `/opt/SUNWnetract/mgmt2.0/docs/api` to find the documentation that identifies the pieces you need to communicate with the `netract` agent.

See the API documentation for:

- `com.sun.ctmgx.MohNames`

- `com.sun.ctmgx.ContainmentTreeMbean`

For JDMK documentation, go to: `/opt/SUNWjdmk/jdmk4.2/1.2/docs`. For an introduction to JDMK, go to `//docs.sun.com`, and search for the *Java Dynamic Management Kit 4.2 Tutorial*.

See the API documentation for:

- `com.sun.jdmk.comm.RmiConnectorAddress`
- `com.sun.jdmk.comm.RmiConnectorClient`

Communicating With the Netra CT Agent

This simple demonstration lets you connect a client with an instance of `netract` agent, beginning in [CODE EXAMPLE 3-2](#). This example represents part one of a three-part example. A detailed explanation follows.

CODE EXAMPLE 3-2 Creating a Client to Communicate With the Netra CT Agent (Part 1)

```
import java.util.Iterator;
import java.util.Set;
import javax.management.ObjectName;
import com.sun.ctmgx.moh.MohNames;
import com.sun.jdmk.ServiceName;
import com.sun.jdmk.comm.RmiConnectorAddress;
import com.sun.jdmk.comm.RmiConnectorClient;

public class Client {

    private RmiConnectorClient connectorClient;
    private RmiConnectorAddress connectorAddress;

    public Client() {
        connectorClient = new RmiConnectorClient();
        connectorAddress = new RmiConnectorAddress();
    }

    public static void main(String[] args) {
        Client client = new Client();
        try {
            client.printContainmentTree();
        } catch (Exception e) {
            e.printStackTrace();
        }
    }
}
```

[CODE EXAMPLE 3-2](#) instantiates the `RmiConnectorClient` and `RmiConnectorAddress`.

The demonstration continues in [CODE EXAMPLE 3-3](#).

Finding the Root Object Name

[CODE EXAMPLE 3-3](#) continues from the previous example. The code example connects to the client and prints the `ContainmentTree` by getting the `ObjectName` of the root `MBean` in the containment hierarchy.

Each `MohNames` instance comes up with `ObjectNames` instances that are accessible via public static fields defined in `MohNames`. This includes the `ContainmentTreeMBean` instance, which provides a mechanism for the user to traverse the containment hierarchy representing the Netra CT system.

CODE EXAMPLE 3-3 Getting the Root MBean Object Name (Part 2)

```
public void printContainmentTree() throws Exception {
    connectorClient.connect(connectorAddress);

    Object[] params = new Object[0];
    String[] signature = new String[0];
    ObjectName rootName =
        (ObjectName)connectorClient.invoke(MohNames.MOH_CONTAINMENT_TREE,
                                           "getRoot", params, signature);
    printSubTree(rootName);

    connectorClient.disconnect();
}
```

This demonstration returns the `ObjectName` of the instance of `NEMBean`. `NEMBean` is the name of the network element `MBean` representing the system as a whole, in other words, the root of the tree.

Now that you have identified the `ObjectName` of the root of the `MOH_CONTAINMENT_TREE` you are ready to traverse the tree and find out what other elements are in the tree.

Traversing the Containment Hierarchy From a Node

Continuing the demonstration from the previous example, in [CODE EXAMPLE 3-4](#) you traverse the `MOH_CONTAINMENT_TREE` from a node, and can get a list of all the nodes on the tree using `getChildren`.

CODE EXAMPLE 3-4 Traversing the Containment Hierarchy From a Node (Part 3)

```
private void printSubTree(ObjectName nodeName) throws Exception {
    System.out.println(nodeName);

    Object[] params = {nodeName};
    String[] signature = {"javax.management.ObjectName"};

    Set children =
        (Set)connectorClient.invoke(MohNames.MOH_CONTAINMENT_TREE,
                                   "getChildren", params, signature);

    for (Iterator itr = children.iterator(); itr.hasNext();) {
        printSubTree((ObjectName)itr.next());
    }
}
```

Here, the `nodeName` is the `ObjectName` of the MBean where the search starts. The line beginning `Set children` gets the children of the specified MBean in the containment hierarchy.

Once you have established the hierarchy of the existing system, your application must receive notification when changes to the system occur. This is the subject of the following section.

Listening for Notifications

This series of examples assumes you continue from the previous three-part example. Return to `/opt/SUNWnetract/mgmt2.0/docs/api` to find documentation.

In the JDMK framework, look at:

- `javax.management.Notification`
- `javax.management.NotificationListener`
- `javax.management.NotificationFilterSupport`
- `javax.management.NotificationFilter`

In MOH documentation you need to look at: `com.sun.ctmgx.moh.MohNames` and `com.sun.ctmgx.moh.Moh.EFDMBean`.

Registering a Notification Listener With EFDMBean Instance

This example continues from the previous examples and shows you how to register a `NotificationListener` using a `NotificationFilter`. You begin by adding a `NotificationListener` that catches communications from the `RmiConnectorClient`.

CODE EXAMPLE 3-5 RMI Example of Listening for MOH Notifications

Registering a `NotificationListener` with a `NotificationFilter`

```
try {
    // accessing MohNames for MOH_DEFAULT_EFD
    //
    connectorClient.addNotificationListener(MohNames.MOH_DEFAULT_EFD, \
                                           aListener, aFilter, null);
}
catch (com.sun.jdmk.comm.CommunicationException ce) {
    try {
        connectorClient.setMode(RmiConnectorClient.PULL_MODE);
        connectorClient.addNotificationListener\
            (MohNames.MOH_DEFAULT_EFD, aListener, aFilter, null);
    }
    catch (Exception e) {
    }
}
```

CODE EXAMPLE 3-5 establishes that `MohNames` can access `MOH_DEFAULT_EFD`. The `EFDMBean` exposes the remote management interface of an event forwarding discriminator managed object.

The `netract` agent of the Netra CT alarm card does not support the `PUSH_MODE`, so the above code will work for any of the `netract` agent instances (those on the host, satellite, and alarm card) in a Netra CT drawer.

Managing Alarms

Before you begin this segment of code example, you should refer back to:

`/opt/SUNWnetract/mgmt2.0/docs/api`

Look at the MOH documentation for:

- `com.sun.ctmgx.moh.AlarmNotification`
- `com.sun.ctmgx.moh.AlarmNotificationFilter`
- `com.sun.ctmgx.moh.AlarmSeverity`
- `com.sun.ctmgx.moh.AlarmSeverityProfileMBean`

■ com.sun.ctmgx.moh.AlarmType

Registering a NotificationListener With an AlarmNotificationFilter

In this section you identify the kinds of alarms the script listens for when events occur. You can specify the level of action; this example listens for critical or major alarms. AlarmNotification represents an alarm notification emitted by an MBean.

CODE EXAMPLE 3-6 Registering a NotificationListener With an AlarmNotificationFilter

```
AlarmNotificationFilter aFilter = new AlarmNotificationFilter();

// interested in all types of alarms
//
aFilter.enableAllAlarmTypes();

// interested in only CRITICAL and MAJOR alarms
//
aFilter.enableSeverity(AlarmSeverity.CRITICAL);
aFilter.enableSeverity(AlarmSeverity.MAJOR);

try {
    connectorClient.addNotificationListener(MohNames.MOH_DEFAULT_EFD, \
                                           aListener, aFilter, null)
}
catch (com.sun.jdmk.comm.CommunicationException ce) {
    connectorClient.setMode(RmiConnectorClient.PULL_MODE);
    connectorClient.addNotificationListener(MohNames.MOH_DEFAULT_EFD, \
                                           aListener, aFilter, null)
}
catch (Exception e) {
}
```

CODE EXAMPLE 3-6 follows the form of the previous example in setting the RmiConnectorClient to PULL_MODE. The alarm filter is set to enableAllAlarmTypes, then refined to enable only AlarmSeverity.CRITICAL and AlarmSeverity.MAJOR.

Using the Default AlarmSeverityProfile

Each netract agent instance comes up with a default instance of `AlarmSeverityProfile` which can be accessed by its object name, `MohNames.MOH_DEFAULT_ASP`. The MBean instances that might generate `AlarmNotifications` will have this default `AlarmSeverityProfile` associated with them. You can associate a new profile any time.

CODE EXAMPLE 3-7 Using the Default AlarmSeverityProfile

```
// Get the alarm severity association of the default profile
//
Object[] allObjs = null;
Object obj = null;
Java.util.Set mySet = null;
Java.util.Map myMap = null;
    try {
        myMap = (Map)connectorClient.invoke(MohNames.MOH_DEFAULT_ASP, \
            "getAlarmSeverityList", null, null);

        mySet = (Set)myMap.keySet();
        allObjs = mySet.toArray();
    } catch(Exception e) {
        e.printStackTrace();
    }

AlarmType aType = null;
AlarmSeverity aSeverity = null;

for (int i = 0; i < mySet.size();i++) {
    try {
        // aType and aSeverity is the association in this
        // default profile
        aType = (AlarmType)allObjs[i];
        aSeverity = (AlarmSeverity)myMap.get(aType);

        // setting the severity of high temp alarm to critical
        //
        if (aType.equals(AlarmType.HIGH_TEMPERATURE)) {
            Object[] params = new Object[2];
            String[] signature = new String[2];
            params[0] = aType;
            params[1] = AlarmSeverity.CRITICAL;
            signature[0] = "com.sun.ctmgx.moh.AlarmType";
            signature[1] = "com.sun.ctmgx.moh.AlarmSeverity";
            connectorClient.invoke(MohNames.MOH_DEFAULT_ASP, \
                "setAlarmSeverity", params, signature);
        }
    }
}
```

CODE EXAMPLE 3-7 Using the Default AlarmSeverityProfile *(Continued)*

```
    } catch(Exception e) {  
        e.printStackTrace();  
    }  
}
```

In [CODE EXAMPLE 3-7](#), the severity level of `HIGH.TEMPERATURE AlarmType` in the default `AlarmSeverityProfile` has been set to `CRITICAL`. The following example shows how to create your own alarm severity profile instances.

Creating Your Own AlarmSeverityProfile

You can create your own `AlarmSeverityProfile` by following [CODE EXAMPLE 3-8](#).

CODE EXAMPLE 3-8 Creating an AlarmSeverityProfile

```
try {  
    // You need to provide the class name to instantiate an MBean,  
    // for AlarmSeverityProfileMBean  
    // the class name string is defined by the constant MohNames.CLASS_NAME_ASP  
    //  
    ObjectName profileName = new ObjectName("NetraCT:name=\  
        AlarmSeverityProfile,id=2");  
    connectorClient.createMBean(MohNames.CLASS_NAME_ASP, profileName, \  
        null,null);  
  
    // To make the profile usable, you need to provide the alarm type and severity  
    // associations  
    //  
    Object[] params = new Object[2];  
    String[] signature = new String[2];  
    signature[0] = "com.sun.ctmgx.moh.AlarmType";  
    signature[1] = "com.sun.ctmgx.moh.AlarmSeverity";  
  
    // For high temperature alarm  
    //  
    params[0] = AlarmType.HIGH_TEMPERATURE;  
    params[1] = AlarmSeverity.CRITICAL;  
    connectorClient.invoke(profileName, \  
        "setAlarmSeverity", params, signature);  
  
    // For high memory utilization alarm  
    //  
    params[0] = AlarmType.HIGH_MEMORY_UTILIZATION;  
    params[1] = AlarmSeverity.MAJOR;
```


CODE EXAMPLE 3-8 Creating an AlarmSeverityProfile (Continued)

```
connectorClient.invoke(profileName, \
    "setAlarmSeverity", params, signature);

// For fan failure alarm (NetraCT agent does not support this alarm
// currently
//
params[0] = AlarmType.FAN_FAILURE;
params[1] = AlarmSeverity.MINOR;
connectorClient.invoke(profileName, \
    "setAlarmSeverity", params, signature);

// For fuse failure alarm (NetraCT agent does not support this alarm
// currently
//
params[0] = AlarmType.FUSE_FAILURE;
params[1] = AlarmSeverity.WARNING;
connectorClient.invoke(profileName, \
    "setAlarmSeverity", params, signature);

} catch (Exception e) {
    e.printStackTrace();
}
```

CODE EXAMPLE 3-8 assigns alarm notifications for high temperature, high memory usage, fan failure, and fuse failure, although the current netract agent does not support alarm notifications for fuse failure. The code example is included here for demonstration purposes.

Assigning a New AlarmSeverityProfile

CODE EXAMPLE 3-9 shows how to assign a new AlarmSeverityProfile to an MBean which can generate AlarmNotifications.

CODE EXAMPLE 3-9 Assigning a New AlarmSeverityProfile

```
try {

    Object[] params = new Object[1];
    String[] signature = new String[1];

    signature[0] = "javax.management.ObjectName";

    // pass the object name of the newly created AlarmSeverityProfileMBean
    // instance
    //
    params[0] = profileName;
```

CODE EXAMPLE 3-9 Assigning a New AlarmSeverityProfile (Continued)

```
// sensorObjectName is the object name of lets say a temperature sensor
// MBean instance
//
connectorClient.invoke(sensorObjectName,\
    "setAlarmSeverityProfilePointer", params, signature);
} catch (Exception e) {
    e.printStackTrace();
}
```

The new AlarmSeverityProfile can be reserved to replace the default profile when required.

Configuring the Agent to Drive Alarm Card Alarm Outputs

The system configuration hierarchy indicates the physical alarm port which corresponds to a termination point, as shown in [“Hardware Resource Hierarchy Showing Managed Object Classes” on page 11](#) and subsequent views. The alarm port termination point supports five alarm interfaces – three for output, two for input. In general, when an alarm occurs the corresponding output alarm pin is driven high based on the alarm severity.

The output alarm pins (alarm0, alarm1, alarm2) are statically mapped into severities of critical, major, and minor respectively.

For example, assume that HIGH_TEMPERATURE is assigned as critical, and HIGH_MEMORY_UTILIZATION is assigned as minor. When a high temperature occurs, alarm0 is driven high to indicate a critical alarm. When a HIGH_MEMORY_UTILIZATION occurs, alarm2 is driven high to indicate a minor alarm.

TABLE 3-2 Example of Alarm Output Mapping

alarm0	alarm1	alarm2
critical	major	minor
HIGH_TEMPERATURE		HIGH_MEMORY_UTILIZATION

In JMX, an alarm is defined as a notification with a severity associated with it. These alarms are assigned as NetworkInterfaceMBeans, each of which represent a network interface object in the system.

You can configure an alarm card agent to drive output alarms from the alarm card on the Netra CT server using MOH as described in the following section.

▼ To Set Up and Use Alarm Features

The following steps show how to configure an agent from the alarm card to correspond with the mapping in [TABLE 3-2](#).

1. Register a notification listener with an `AlarmNotificationFilter`.

Use the examples beginning [“Registering a NotificationListener With an AlarmNotificationFilter” on page 30](#), and modify the default to listen for critical or major alarms. Return to the start of this chapter for help in getting an `ObjectName`.

2. Develop an `AlarmSeverityProfile` based on the default profile.

An `AlarmSeverityProfile` (ASP) contains multiple entries, and can be assigned to several alarm-generating objects. Some entries in the profile might not be used by an object, because that object might not be generating that specific kind of alarm. The default instance of `AlarmSeverityProfile` can be accessed by its object name, `MohNames.MOH_DEFAULT_ASP`.

3. Assign the `AlarmSeverityProfile` to the corresponding objects.

- Assign `HIGH_TEMPERATURE` to the corresponding CPU thermistor `sensorObjectName`.
- Assign `HIGH_MEMORY_UTILIZATION` to the corresponding `CpuCardEquipment` object.

In [CODE EXAMPLE 3-10](#) extracted from [“Using the Default AlarmSeverityProfile” on page 31](#), the severity level of `HIGH_TEMPERATURE` alarm type in the default ASP has been set to `CRITICAL` corresponding with `alarm0`.

CODE EXAMPLE 3-10 Extract of Using the Default Alarm Severity Profile

```
// Get the alarm severity association of the default profile
//
<snip>
.....

        // setting the severity of high temp alarm to critical
        //
        if (aType.equals(AlarmType.HIGH_TEMPERATURE)) {
            Object[] params = new Object[2];
            String[] signature = new String[2];
            params[0] = aType;
            params[1] = AlarmSeverity.CRITICAL;
            signature[0] = "com.sun.ctmgx.moh.AlarmType";
            signature[1] = "com.sun.ctmgx.moh.AlarmSeverity";
            connectorClient.invoke(MohNames.MOH_DEFAULT_ASP, \
                "setAlarmSeverity", params, signature);
```

CODE EXAMPLE 3-10 Extract of Using the Default Alarm Severity Profile (*Continued*)

```
.....  
<unsnip>
```

Any number of objects are capable of generating an alarm. If you assign this profile to a particular object, whenever a hardware failure of that object occurs, the netrtract agent refers to the profile and responds as you have specified.

[CODE EXAMPLE 3-11](#) creates your own `AlarmSeverityProfile` instances based on these examples. In this case, the `sensorObjectName` is the object name of a temperature sensor MBean instance.

CODE EXAMPLE 3-11 Extract of Assigning a New `AlarmSeverityProfile`

```
try {  
    Object[] params = new Object[1];  
    String[] signature = new String[1];  
  
    signature[0] = "javax.management.ObjectName";  
  
    // pass the object name of the newly created AlarmSeverityProfileMBean  
    // instance  
    //  
    params[0] = profileName;  
  
    // sensorObjectName is the object name of lets say a temperature sensor  
    // MBean instance  
    //  
    connectorClient.invoke(sensorObjectName,\  
        "setAlarmSeverityProfilePointer", params, signature);  
} catch (Exception e) {  
    e.printStackTrace();  
}
```

The new alarm severity profile replaces the default profile when required.

You can create several alarm severity profiles, each specifying a different response. One might designate fan failure as critical, another might designate high temperature as major. You then assign the appropriate profile to the object.

Clearing Alarms

Alarms are cleared automatically when each alarm relay is driven low. OperationalState will accordingly be shown to be enabled, disabled, or unknown.

Software Monitoring

The following code examples help monitor software events. The series begins in [CODE EXAMPLE 3-12](#) with establishing the printService to print system status reports, then gathers the list of software services and their associated daemons.

CODE EXAMPLE 3-12 Software Monitor Test (Part 1)

```
private void printService(ObjectName objName) {
    try {
        String name = (String)connectorClient.getAttribute(objName, "Name");
        String status = (String)connectorClient.getAttribute(objName, "Status");
        Integer polling_interval = \
            (Integer)connectorClient.getAttribute(objName, "PollingInterval");
        System.out.println("Name: " + name);
        System.out.println("Status: " + status);
        System.out.println("Polling Interval: " + polling_interval);
    }catch(Exception e) {
        e.printStackTrace();
    }
}

private void printDaemon(ObjectName objName) {
    try {
        String name = (String)connectorClient.getAttribute(objName, "Name");
        String status = (String)connectorClient.getAttribute(objName, "State");
        Integer retry = (Integer)connectorClient.getAttribute(objName, \
            "CurrentRetryCount");
        Integer maxretry = (Integer)connectorClient.getAttribute(objName, \
            "MaxRetryCount");

        System.out.println("name: " + name);
        System.out.println("state: " + status);
        System.out.println("retry: " + retry);
        System.out.println("maxretry: " + maxretry);
    }catch(Exception e) {e.printStackTrace();}
}

// This method traverses through the hierarchy of software monitor and
// prints out all the software services and the daemons.
```

CODE EXAMPLE 3-12 Software Monitor Test (Part 1)

```
private void test() {
    Object[] allObjs = null;
    Set swServiceList = null;
    ObjectName myObjName = null;
    try {
// Get the list of all software services */
swServiceList =(Set)connectorClient.invoke\
(MohNames.MOH_SOFTWARE_MONITOR, "getSoftwareServiceList", null, null);
allObjs = swServiceList.toArray();
    }
```

CODE EXAMPLE 3-12 builds on previous examples to establish the status of `connectorClient`, and examine the hierarchy of the `swServiceList` in order to find existing services and running daemons.

The following segment of code collects the attributes of each software service so that the service can be registered to receive event notification.

CODE EXAMPLE 3-13 Traversing the Software Service List (Part 2)

```
//Traverse through the software service list and print out the attributes
//of each software service
for (int i = 0; i < swServiceList.size();i++) {
myObjName = (ObjectName)allObjs[i];
System.out.println("service : "+ ((ObjectName)allObjs[i]).toString());
printService(myObjName);
// Register the software service to receive the event notifications
connectorClient.addNotificationListener(myObjName, this, null, null);
}
```

CODE EXAMPLE 3-13 traverses `swServiceList` and adds `NotificationListener` to the `connectorClient`.

The final code segment gets the list of daemons that support the service, and prints out the daemon attributes for event notification.

CODE EXAMPLE 3-14 Getting the List of Service Daemons (Part 3)

```
// For each service, get the list of daemons that support the service
ObjectName[] daemonList =\
(ObjectName[])connectorClient.getAttribute(myObjName, "DaemonList");
if (daemonList != null && daemonList.length > 0){
System.out.print("Daemon List: ");
Integer retry = null;
Integer maxretry = null;
// For each daemon, print out all attributes of the daemon.
for (int k= 0;k < daemonList.length;k++) {
printDaemon(daemonList[k]);
}
```

CODE EXAMPLE 3-14 Getting the List of Service Daemons (Part 3)

```
        // register the daemon to receive the event notifications
        connectorClient.addNotificationListener\
            (daemonList[k],this,null,null);
    }
}
} catch(Exception e) {throw new UncheckedException(e);}
}
```

CODE EXAMPLE 3-14 establishes a `DaemonList` for each service and prints out the attributes of each daemon. Finally, the code registers these daemons to receive notice of events with `addNotificationListener`.

For further information, look at `/opt/SUNWnetract/mgmt2.0/docs/api` which details all the MOH interfaces and classes that are provided for the Netra CT system software.

Netra CT Element Management Agent API

This chapter contains the application programming interfaces (API) of the Netra CT element management agent software and includes the following sections:

- [“Interface Overview” on page 42](#)
- [“How the API Sections are Organized” on page 44](#)
- [“Netra CT Management Agent Interfaces and Classes” on page 45](#)

Interface Overview

Netra CT management agent uses the Java Dynamic Management Kit (JDMK) framework as a Java API which provides the management capability for the Netra CT system.

JDMK supports JMX, which is a standard set of APIs for network and client management. JDMK provides an extended API along with different communication protocol adapters such as Remote Method Invocation (RMI), HTTP, HTML, and Simple Network Management Protocol (SNMP).

These protocol adapters are used to communicate with instances of JDMK agents; Netra CT management agent supports SNMP and RMI communication protocols.

You can find an introduction to the JDMK, tutorials, code samples, and APIs on the Sun Developer Network web site: <http://java.sun.com>

Summary of JDMK

JDMK's API and development tools can help you develop distributed management applications. The JDMK allows resources of one host to be monitored from another host.

A resource can be any entity, physical or virtual, that you want to monitor through your network. Physical resources include network elements, and virtual resources include applications operating on a host. A resource can be seen through its management interface, where its attributes, operations, and notifications are accessible by a management agent.

In order for a management agent to monitor a resource, the resource must be developed as a managed bean (MBean), which is Java object that represents the resource's management interface. If the resource itself is a Java application, it can be its own MBean. Otherwise, an MBean is a Java representation of a device.

In the JDMK model, a Java Dynamic Management agent follows the client-server model, in which an agent responds to the management requests from any number of client applications that wish to access its resources. The central component of an agent is the MBean server, which is a registry for MBean instances and provides the framework that allows agent services to interact with MBeans.

The JDMK provides protocol connector interfaces that allow remote applications to access agent applications and their resources. Remote method invocation (RMI) and HTTP are two such JDMK supported protocols that enable a Java client application running on one system to access the resources and methods of another Java server application running on a different system.

FIGURE 4-1 displays the location of the RMI/HTTP protocols between an agent application and a remote manager application.




FIGURE 4-1 Key Components of the Java Dynamic Management Kit

In **FIGURE 4-1**, a resource and an agent service are registered as MBeans with the agent application's MBean server. The application agent also contains a connector server for the RMI/HTTP protocols. The remote manager application is a Java application running on a distant host system. The manager contains the RMI/HTTP connector client and proxy MBeans representing the resource and service. When the RMI/HTTP connector client establishes the connection with the agent's RMI/HTTP connector server, the other components of the application can issue management requests to the agent.

Typically, you would first determine the management interface of your resource, that is, the information needed to manage it. This information is expressed as attributes and operations. An attribute is a value of any type that a manager can get or set remotely. An operation is a method with any signature and any return type that the manager can invoke remotely.

As specified by the Java Management extensions for instrumentation, all attributes and operations are explicitly listed in an MBean interface. This Java interface defines the full management interface of an MBean. The interface must have the same name as the class that implements it, followed by the MBean suffix. Since the interface and its implementation are usually in different files, two files make up a standard

MBean. For example, the management interface of the class `SimpleStandard` (in the file `SimpleStandard.java`) is defined in the interface `SimpleStandardMBean` (in the file `SimpleStandardMBean.java`).

For a complete discussion of JDMK components and protocols, refer to the Java Dynamic Management Kit documentation set found on the Solaris documentation website, <http://docs.sun.com>. For additional information of JDMK and the RMI/HTTP protocol, refer to the documentation, tutorials, code samples, and APIs found on the Java Developers website: <http://developer.java.sun.com>.

Viewing the Netra CT Management Agent API Online

The entire Netra CT RMI API specification can be viewed online as cross-referenced HTML pages. By default, these HTML pages are installed in the following directory:

```
/opt/SUNWnetract/mgmt2.0/docs/api/com/sun/ctmgx/moh
```

You can view an index of all of these pages by opening the following link in a web browser:

```
file:///opt/SUNWnetract/mgmt2.0/docs/api/index.html
```

You can view additional Java API specification on the java.sun.com webpage at:

```
http://java.sun.com/apis.html
```

How the API Sections are Organized

The following sections in this chapter list the classes of the Netra CT RMI application programming interface.

Each class, interface, inner class, and inner interface has its own separate section. Each of these sections have three subsections consisting of a class or interface description, summary tables, and detailed member descriptions of the following:

- Class inheritance diagram
- Direct subclasses
- All known subinterfaces
- All known implementing classes
- Class or interface declaration
- Class or interface description
- Inner class summary

- Field summary
- Constructor summary
- Method summary
- Field detail
- Constructor detail
- Method detail

Each summary entry contains the first sentence from the detailed description for that item. The summary entries are alphabetical, while the detailed descriptions are in the order they appear in the source code. This preserves the logical groupings established by the programmer.

Netra CT Management Agent Interfaces and Classes

[TABLE 4-1](#) lists the management agent interfaces and [TABLE 4-2](#) lists the management agent classes included in the Netra CT RMI API. In these tables, the term *expose* refers to the encapsulation of the object’s variables inside a nucleus. This encapsulation allows for exposing (allowing access to) or hiding (denying access to) an object’s access methods, which provides for greater modularity.

Detailed descriptions of the interfaces reside in the `/opt/SUNWnetract/mgmt2.0/docs/api` directory.

TABLE 4-1 Netra CT Management Agent Interfaces

Interface	Description
AlarmCardPluginMBean	Describes the management interface of the AlarmCardPluginMBean
AlarmSeverityProfileMBean	Describes the management interface of the AlarmSeverityProfileMBean.
CgtpServiceMBean	Describes the management interface of the CgtpServiceMBean.
ContainmentTreeMBean	Describes the management interface of the ContainmentTreeMBean.
CpiSlotMBean	Describes the management interface of the CpiSlotMBean objects.
CpuCardEquipmentMBean	Describes the management interface for the CpuCardEquipmentMBean.

TABLE 4-1 Netra CT Management Agent Interfaces (*Continued*)

Interface	Description
CpuPluginMBean	Describes the management interface of the CpuPluginMBean board objects as perceived from the alarm card MOH.
DaemonMBean	Describes the interface for the DaemonMBean.
EFDMBean	Describes the management interface of the EFDMBean.
EquipmentHolderMBean	Describes the management interface of the EquipmentHolderMBean.
EquipmentMBean	Describes the interface of the EquipmentMBean.
EtherIfStatsMBean	Describes the management interface of the EtherIfStatsMBean I/O Statistics Monitoring service.
FullLogMBean	Describes the interface of the FullLogMBean
IpServiceMBean	Describes the interface of IpServiceMBean, the UNIX File System (UFS) service.
LOLMBean	Describes the management interface of LOLMBean, the Latest Occurrence Log MBean
NEMBean	Describes the management interface of the NEMBean.
NetworkInterfaceMBean	Describes the management interface for NetworkInterfaceMBean.
NfsServiceMBean	Describes the management interface of NfsServiceMBean, the Network File System (NFS) Monitor service.
NumericSensorMBean	Describes the interface for NumericSensorMBean.
PlugInUnitMBean	Describes the management interface of the PlugInUnitMBean.
RnfsServiceMBean	Describes the interface of RnfsServiceMBean, the Reliable Network File System (RNFS) Monitor service.
SensorMBean	Describes the interface for the SensorMBean.
SlotMBean	Describes the management interface of the SlotMBean.
SoftwareMonitorMBean	Describes the interface of the SoftwareMonitorMBean.
SoftwareServiceMBean	Describes the interface of SoftwareServiceMBean.
TcpServiceMBean	Describes the interface of TcpServiceMBean, the Transmission Control Protocol (TCP) service.

TABLE 4-1 Netra CT Management Agent Interfaces (*Continued*)

Interface	Description
TerminationPointMBean	Describes the management interface of the TerminationPointMBean.
UdpServiceMBean	Describes the interface of UdpServiceMBean, the User Datagram Protocol (UDP) service.
UfsServiceMBean	Describes the interface of UfsServiceMBean, the UNIX File System (UFS) service.

TABLE 4-2 Netra CT Management Agent Classes

Class	Description
AdministrativeState	Defines the AdministrativeState of the device.
AlarmNotification	The AlarmNotification class represents an alarm notification emitted by an MBean.
AlarmNotificationFilter	Allows you to filter AlarmNotification notifications by selecting the types and severities of interest.
AlarmSeverity	Defines the AlarmSeverity objects for use with AlarmNotification.
AlarmType	This class is an enumeration of predefined AlarmType, user need to use one of the predefined types to construct an AlarmNotification object.
AttributeChangeNotification	Provides definitions of the AttributeChangeNotification sent by MBeans.
AttributeChangeNotificationFilter	The filtering of the AttributeChangeNotificationFilter is performed on the name of the observed attribute.
AuthClient	AuthClient class defines the client utility routines, particularly for authentication.
AvailabilityStatus	Defines the AvailabilityStatus of the plug-in unit object.
EquipmentHolderType	Describes the management interface of the EquipmentHolderType.
LogFullAction	Describes the action to perform when LogFullAction indicates the log is full.

TABLE 4-2 Netra CT Management Agent Classes (*Continued*)

Class	Description
MohNames	Defines MohNames, the public constants or static variables for MOH user to communicate to the MBean server.
ObjectCreationNotification	Defines ObjectCreationNotification, the creation notifications sent by MBeans.
ObjectDeletionNotification	Defines ObjectDeletionNotification, the deletion notifications sent by MBeans.
OperationalState	Defines the OperationalState of a device, equipment, or plug-in.
SlotStatus	Defines SlotStatus, the status of the slot object.
StateChangeNotification	Defines StateChangeNotification, the state change notifications sent by MBeans.
StateChangeNotificationFilter	Describes StateChangeNotificationFilter, the filtering performed on the name of the observed attribute.

Simple Network Management Protocol

This chapter describes the Netra CT server Simple Network Management Protocol (SNMP) support, and provides a useful example. This chapter contains the following sections:

- [“SNMP Overview” on page 50](#)
- [“Netra CT System SNMP Representation” on page 52](#)
- [“ENTITY-MIB” on page 52](#)
- [“SUN-SNMP-NETRA-CT-MIB” on page 55](#)
- [“Changing Midplane FRU-ID” on page 66](#)

SNMP Overview

The most widespread legacy architecture for network and device management is SNMP, for which the Java DMK provides a complete toolkit. This gives you the advantages of developing both Java Dynamic Management agents and managers that are interoperable with existing management systems.

SNMP network protocol enables devices to be managed remotely by a Network Management Station (NMS). To be managed, a device must have an SNMP agent associated with it. The agent receives requests for data representing the state of the device and provides an appropriate response. The agent can also control the state of the device. Additionally, the agent can generate SNMP traps, which are unsolicited messages sent to selected NMSs to signal significant events relating to the device.

The Sun Netra SNMP Management Agent is an intelligent SNMP v2 agent for continuously monitoring key hardware variables. You can generate and collect value-add reports collected by remote monitoring. Using Sun Netra SNMP Management Agent's generic management interface and comprehensive event mechanisms, you can dynamically build configuration and health status data, thus reducing development costs.

Management Information Base

To manage and monitor devices, the characteristics of the devices must be represented using a format known to both the agent and the NMS. These characteristics can represent physical properties such as fan speeds, or services such as routing tables. The data structure defining these characteristics is known as a Management Information Base (MIB). This data model is typically organized into tables, but can also include simple values. An example of the former is routing tables, and an example of the latter is a timestamp indicating the time at which the agent was started.

A MIB is a text file, written in abstract syntax notation one (ASN.1) notation, which describes the variables containing the information that SNMP can access. The variables described in a MIB, which are also called MIB objects, are the items that can be monitored using SNMP. There is one MIB object for each element being monitored. All MIBs are, in fact, part of one large hierarchical structure, with leaf nodes containing unique identifiers, data types, and access rights for each variable and the paths providing classifications. A standard path structure includes branches for private subtrees.

For reference, the structure of the MIBs for SNMPv2 is defined by its Structure of Management Information (SMI) defined in the RFC2578 document. This SMI defines the syntax and basic data types available to MIBs. The Textual Conventions (type definitions) defined in the RFC2579 document define additional data types and enumerations.

Before an NMS can manage a device through its agent, the MIB corresponding to the data presented by the agent must be loaded into the NMS. The mechanism for doing this varies depending on the implementation of the network management software. This gives the NMS the information required to address and correctly interpret the data model presented by the agent. Note that MIBs can reference definitions in other MIBs, so to use a given MIB, it might be necessary to load others.

Object Identifiers

The MIB defines a virtual datastore accessible by way of the SNMP software, the content being provided either by corresponding data maintained by the agent, or by the agent obtaining the required data on demand from the managed device. For writes of data by the NMS to this virtual data, the agent typically performs some action affecting the state either of itself or the managed device.

To address the content of this virtual datastore, the MIB is defined in terms of object identifiers (OIDs) which uniquely identify each data entry. An OID consists of an hierarchically arranged sequence of integers providing a unique name space. Each assigned integer has a associated text name. For example, the OID `1.3.6.1` corresponds to the OID `iso.org.dod.internet` and `1.3.6.1.4` corresponds to the OID `iso.org.dod.internet.private`. The numeric form is used within SNMP protocol transactions, whereas the text form is used in user interfaces to aid readability. Objects represented by such OIDs are commonly referred to by the last component of their name as a shorthand form. To avoid confusion arising from this convention, it is normal to apply a MIB-specific prefix, such as `netract`, to all object names defined therein.

All addressable objects defined in the MIB have associated maximum access rights (for instance, read-only or read-write), which determine what operations the NMS permits the operator to attempt. The agent can limit access rights as required; that is, it is able to refuse writes to objects that are considered read-write. This refusal can be done on the grounds of applicability of the operation to the object being addressed, or on the basis of security restrictions that can limit certain operations to restricted sets of NMS. The mechanism used to communicate security access rights is *community strings*. These text strings, such as `private` and `public`, are passed with each SNMP data request.

Much of the data content defined by MIBs is of a tabular form, organized as entries consisting of a sequence of objects (each with their own OIDs). For example, a table of fan characteristics could consist of a number of rows, one per fan, with each row

containing columns corresponding to the current speed, the expected speed, and the minimum acceptable speed. The addressing of the rows within the table can be a simple single dimensional index (a row number within the table, for example, 6), or a more complex, multidimensional, instance specifier such as an IP address and port number (for example, 127.0.0.1, 1234). In either case, a specific data item within a table is addressed by specifying the OID giving its prefix (for example, myFanTable.myFanEntry.myCurrentFanSpeed) with a suffix instance specifier (for example, 127.0.0.1.1234 from the previous example) to give myFanTable.myFanEntry.myCurrentFanSpeed.127.0.0.1.1234.

Each table definition within the MIB has an INDEX clause that defines which instance specifiers to use to select a given entry. The SMI defining the MIB syntax provides an important capability whereby tables can be extended to add additional entries, effectively adding extra columns to the table. This is achieved by defining a table with an INDEX clause that is a duplicate of that of the table being extended.

Netra CT System SNMP Representation

The Netra CT software uses these SNMP MIBs to present the network information model:

- ENTITY-MIB (RFC 2037)
- IF-MIB (RFC 2863)
- HOST-RESOURCES-MIB (RFC 2790)
- SUN-SNMP-NETRA-CT-MIB

ENTITY-MIB

The ENTITY-MIB is defined by the IETF standard RFC2037. The ENTITY-MIB provides a mechanism for presenting hierarchies of physical entities using SNMP tables.

The Netra CT information model uses the ENTITY-MIB to provide:

- A hierarchy of hardware resources – relationships between managed objects
- Common hardware resource characteristics – a mapping of common attributes from the GNIM Top, Equipment, and Termination Point classes

This information is presented using SNMP tables:

- Physical Entity Table (entPhysicalTable)

This table contains one row per hardware resource. These rows are called *entries*, and a particular row is referred to as an *instance*. Each entry contains the physical class (entPhysicalClass) and common characteristics of the hardware resource. Each entry has a unique index (entPhysicalIndex) and contains a reference (entPhysicalContainedIn) that points to the row of the hardware resource which acts as the *container* for this resource.

FIGURE 5-1 and TABLE 5-1 show how an example hierarchy of hardware resources are presented using the ENTITY-MIB.

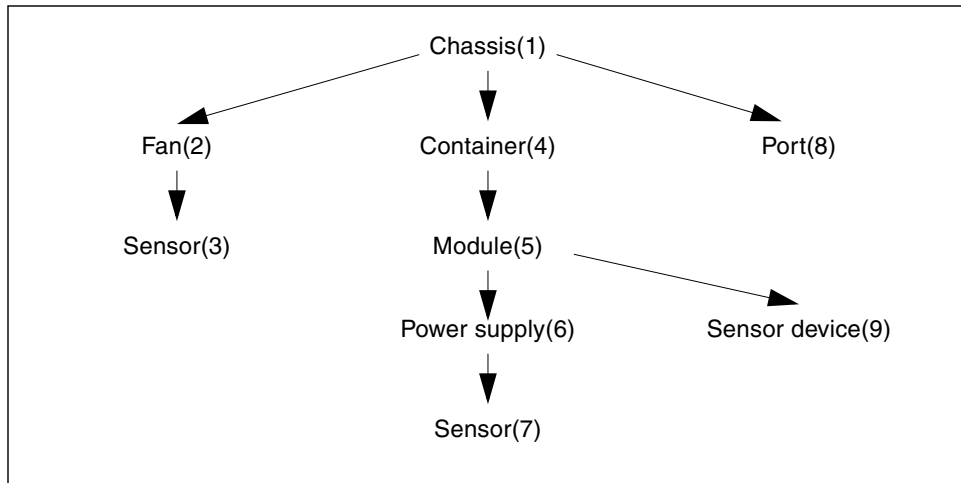


FIGURE 5-1 Hardware Resource Hierarchy

TABLE 5-1 Physical Entity Table

entPhysicalIndex	entPhysicalClass	entPhysicalContainedIn	...
1	chassis	0	...
2	fan	1	...
3	sensor	2	...
4	container	1	...
5	module	4	...
6	power supply	5	...
7	sensor	6	...

TABLE 5-1 Physical Entity Table

<code>entPhysicalIndex</code>	<code>entPhysicalClass</code>	<code>entPhysicalContainedIn</code>	...
8	port	1	...
9	other	5	...
10	other	5	...

The Netra CT Management Agent uses values for `entPhysicalIndex` and `ifIndex` that might not be contiguous, but are within the range of permitted values.

IF-MIB

The IF-MIB is defined by the IETF standard RFC 2863. The IF-MIB provides information about the network interfaces of the server. The information is presented using the `ifTable`. The `ifTable` contains a row for each network interface. The `ifTable` includes columns which describe the interface (`ifDescr`), indicate the type of interface (`ifType`), and the indicate the status of the interface (`ifOperStatus`).

HOST-RESOURCES-MIB

The `Host-Resources-MIB` is defined in RFC 2790.

Host Resources Running Software Table

The Host Resources Running Software Table (`hrSWRunTable`) contains information about the software that is running on the network element (for example, NFS, TFTP, and CGTP). When an application or daemon under the monitor is running, the MOH Software Module adds an entry into the `hrSWRunTable` and will send to the client the `netraCtRunningSwCreated` trap. When an application or a daemon stops running, the MOH Software Module sends the `netraCtRunningSwChanged` trap with `hrSWRunStatus` is invalid. The MOH Software Module only deletes the entry from the `hrSWRunTable` and sends the `netraCtRunningSwDeleted` trap when the service is uninstalled from the system.

Host Resources Installed Software Table

The Host Resources Installed Software Table (`hrSWInstalledTable`) contains information about the software installed on the network element (for example, installation packages related to NFS, CGTP, and so on). `netraCtInstalledSwCreated`, `netraCtInstalledSwDeleted` and `netraCtInstalledSwChanged` are traps sent to the client corresponding to the software package installed event, software package uninstalled event, and different version of the existing software package installed event.

SUN-SNMP-NETRA-CT-MIB

This section describes the `SUN-SNMP-NETRA-CT-MIB`, which is the SNMP version of the Netra CT network element view.

A brief description of each of the groups that comprise the MIB module is provided in the following subsections:

- [“Netra CT Network Element High-Level Objects” on page 56](#)
- [“Physical Path Termination Point Table” on page 56](#)
- [“Equipment Table” on page 57](#)
- [“Plug-in Unit Table” on page 59](#)
- [“Hardware Unit to Running Software Relationship Table” on page 60](#)
- [“Hardware Unit to Installed Software Relationship Table” on page 60](#)
- [“Alarm Severity Identifier Textual Convention” on page 61](#)
- [“Alarm Severity Profile Table” on page 61](#)
- [“Alarm Severity Table” on page 61](#)
- [“Trap Forwarding Table” on page 62](#)
- [“MIB Notification Types” on page 63](#)

For more information, refer to the MIB file that is available as part of the software package at the default location:

```
/opt/SUNWnetract/mgmt2.0/mibs/SUN-SNMP-NETRA-CT-MIB.mib
```

Netra CT Network Element High-Level Objects

The `SUN-SNMP-NETRA-CT-MIB` module representation of high-level objects in the Netra CT network element (NE) is composed of the elements in [TABLE 5-2](#):

TABLE 5-2 `SUN-SNMP-NETRA-CT-MIB` Netra CT NE High-Level Objects

Field	Description
Vendor	The vendor of the Netra CT network element.
Version	The version of the Netra CT network element.
Start Time	The time at which the agent was last started; in other words, the time at which <code>sysUpTime</code> was zero (0).
Alarm Severity Index	An index into the alarm severity profile table, specifying the severity assignments for Netra CT alarms reported for the Netra CT network element. The default value for this object is zero (0).
Suppress Zero Stats	When the value of this object is true, no entry will be created in any of the historical statistics tables for intervals in which all counts are zero. The default value for this object is true (1).

Physical Path Termination Point Table

The Netra CT Physical Path Termination Point Table extends the `entPhysicalTable`. Each entry of this table represents a Physical Path Termination Point within the Netra CT NE. The `SUN-SNMP-NETRA-CT-MIB` module representation of a physical path termination point is composed of the elements shown in [TABLE 5-3](#):

TABLE 5-3 `SUN-SNMP-NETRA-CT-MIB` Physical Path Termination Point Table

Field	Description
Physical Path Termination Point Hardware Unit Index	Specifies the index of the entry in the <code>entPhysicalTable</code> that represents the device (that is, a card) on which the physical path terminates.

TABLE 5-3 SUN-SNMP-NETRA-CT-MIB Physical Path Termination Point Table

Field	Description
Physical Path Termination Point Port ID	Identifies the port within the card identified by the hardware unit index on which the physical path terminates.
Physical Path Termination Point Port Label	Provides the external label string for the physical path termination point entry. If there is no label, the value is a zero-length display string.
Physical Path Termination Point Port Alarm Severity Index	Specifies the index of the entry in the communications alarm severity profile table that should be used. The default value of this object is zero (0).

Equipment Table

The Netra CT Equipment Table extends the entPhysicalTable. Each entry in this table represents a piece of equipment within the Netra CT NE that neither is nor accepts a replaceable plug-in unit. The SUN-SNMP-NETRA-CT-MIB module representation of an equipment is composed of the elements shown in [TABLE 5-4](#):

TABLE 5-4 SUN-SNMP-NETRA-CT-MIB Equipment Table

Field	Description
Equipment Administration Status	Used by the administrator to lock and unlock the object.
Equipment Location	The specific or general location of the component.
Equipment Operating Status	Identifies whether or not the component is capable of performing its normal functions.
Equipment Vendor	The vendor of the component.
Equipment Version	The version of the component.
Equipment User Label	A user-friendly name for the piece of equipment. The default value of this object is the null string.
Equipment Alarm Severity Index	An index into the alarm severity profile table, specifying the severity assignments for Netra CT alarms reported for this component. The default value of this object is zero (0).

Equipment Holder Table

The Netra CT Equipment Holder table extends the `entPhysicalTable`. Each entry in this table represents a component within the Netra CT NE that accepts a replaceable plug-in unit. The `SUN-SNMP-NETRA-CT-MIB` module representation of an equipment holder is composed of the elements shown in [TABLE 5-5](#):

TABLE 5-5 SUN-SNMP-NETRA-CT-MIB Equipment Holder Table

Field	Description
Equipment Holder Type	The type of the component.
Equipment Holder Acceptable Types	The types of plug-in units that can be supported by the slot, separated by newline characters. This attribute is present only when the Equipment Holder represents a slot.
Equipment Holder Slot Status	Identifies whether or not a plug-in unit is present in the slot. This attribute is present only when the Equipment Holder represents a slot.
Equipment Holder Label	Provides the external label string for the holder entry. If there is no label, the value is a zero-length display string.
Equipment Holder Software Load	An index into the installed software table, specifying the software that is to be loaded into the plug-in unit whenever an automatic reload of software is needed. This attribute is present only when the Equipment Holder represents a slot.
Equipment Holder Plug-In Unit Acceptable	This field is true when the plug-in unit contained in the equipment holder is supported.
Equipment Holder Alarm Severity Index	An index into the alarm severity profile table, specifying the severity assignments for Netra CT alarms reported for this component.

Plug-in Unit Table

The Plug-In Unit Table extends the `entPhysicalTable`. Each entry of this table represents a piece of equipment within the Netra CT NE that is inserted into and removed from an Equipment Holder. The `SUN-SNMP-NETRA-CT-MIB` module representation of a plug-in unit is composed of the elements shown in [TABLE 5-6](#).

TABLE 5-6 SUN-SNMP-NETRA-CT-MIB Plug-In Unit Table

Field	Description
Plug-In Unit Administration Status	Used by the administrator to lock and unlock the object. Values are: up (1) and down (2).
Plug-In Unit Availability Status	Provides further information regarding the state of the component. Value are: available (1), inTest (2), failed (3), powerOff (4), notInstalled (5), offline (6), dependency (7), and unknown (8).
Plug-In Unit Operative Status	Identifies whether or not the component is capable of performing its normal functions. Values are: up (1), down (2), and unknown (3).
Plug-In Unit Vendor	The vendor of the component.
Plug-In Unit Version	The version of the component.
Plug-In Unit Label	Provides the external label string for the plug-in entry. If there is no label, the value is a zero-length display string.
Plug-In Unit Alarm Severity Index	An index into the alarm severity profile table, specifying the severity assignments for Netra CT alarms reported for this component. The default value of this object is zero (0).

Hardware Unit to Running Software Relationship Table

The Netra CT Hardware Unit to Running Software Relationship Table describes the software that is running on each hardware unit in the Netra CT NE. Each entry of this table identifies an entry in the `entPhysicalTable` and one in the `hrSWInstalledTable`.

The `SUN-SNMP-NETRA-CT-MIB` hardware unit to running software relationship table is composed of the elements shown in [TABLE 5-7](#).

TABLE 5-7 SUN-SNMP-NETRA-CT-MIB Hardware Unit to Running Software Relation Table

Field	Description
Hardware Running Software to Hardware Index	The index, in the <code>entPhysicalTable</code> , of the containing hardware unit in this pair.
Hardware Running Software Index	A unique number within the context of the containing hardware unit.
Hardware Running Software to Software Index	An index into the Netra CT Hardware Unit to Running Software relationship table.

Hardware Unit to Installed Software Relationship Table

The Netra CT Hardware Unit to Install Software Relationship Table describes the software that is installed on each hardware unit in the Netra CT NE. Each entry of this table identifies an entry in the `entPhysicalTable` and one in the `hrSWInstalledTable`. The `SUN-SNMP-NETRA-CT-MIB` hardware unit to installed software relationship table is composed of the elements shown in [TABLE 5-8](#).

TABLE 5-8 SUN-SNMP-NETRA-CT-MIB Hardware Unit to Installed Software Relationship Table

Field	Description
Hardware Installed Software to Hardware Index	The index, in the <code>entPhysicalTable</code> , of the containing physical entity in this pair.
Hardware Installed Software Index	A unique number within the context of the containing hardware unit.
Hardware Installed Software to Software Index	The index, in the <code>hrSWInstalledTable</code> , of the software product represented by this entry.
Hardware to Software Alarm Severity Index	An index into the alarm severity profile table, specifying the severity assignments for Netra CT alarms reported for this piece of software installed on the hardware unit. The default value of this object is zero.
Hardware Installed Software to Hardware Index	The index, in the <code>entPhysicalTable</code> , of the containing physical entity in this pair.

Alarm Severity Identifier Textual Convention

The SUN-SNMP-NETRA-CT-MIB alarm severity identifier textual conventions consist of the elements shown in [TABLE 5-9](#).

TABLE 5-9 SUN-SNMP-NETRA-CT-MIB Alarm Severity Identifier Textual Conventions

Field	Description
Alarm Log Severity	The value of this object identifies the severity of an alarm in the log. Values are: cleared (-1), indeterminate (0), critical (1), major (2), minor (3), and warning (4).
Alarm Severity	The value of this object identifies the severity of an alarm that has occurred. Values are: indeterminate (0), critical (1), major (2), minor (3), and warning (4). (Note that there is no value corresponding to 'cleared'.)

Alarm Severity Profile Table

The Netra CT alarm severity profile table specifies which profiles exist. Creating or deleting an entry in this table automatically creates or deletes the corresponding entries in the `netraCtAlarmSeverityTable`. Each entry of this table represents a group of severities, one for each alarm type in the communications alarm group. The SUN-SNMP-NETRA-CT-MIB alarm severity profile table consists of the elements shown in [TABLE 5-10](#).

TABLE 5-10 SUN-SNMP-NETRA-CT-MIB Alarm Severity Profile Table

Field	Description
Alarm Severity Profile Index	A number identifying this alarm severity profile.
Alarm Severity Profile Row Status	This object is used to create a new row or to delete an existing row in the table.

Alarm Severity Table

The Netra CT alarm severity table associates profile index and trap ID pairs with severities to be used for Netra CT alarm traps that have occurred. (Note that this table does not apply to cleared alarms). An entry in this table associates an alarm severity profile index and trap ID pair with a severity. Deleting a particular profile's row in the alarm severity profile table deletes all rows in this table with the same profile index. Conceptually, rows corresponding to all possible trap IDs are created

in this table when a new alarm severity profile is created, but the agent returns a default value except for those few traps for which values have been set. The alarm severity table elements are listed in [TABLE 5-11](#).

TABLE 5-11 SUN-SNMP-NETRA-CT-MIB Alarm Severity Table

Field	Description
Alarm Severity Trap ID	The ID of the trap type to which this entry applies.

Trap Forwarding Table

The Netra CT Trap forwarding discriminator table specifies which traps will be sent to which management system. Each entry of this table contents information about a group of traps to be sent to a particular IP address. This is used as the value of the object `netraCtForwardedTrapObject` when traps from all objects are to be forwarded, or when there is only one object of the type that forwards the specified trap type. The elements for this table are shown in [TABLE 5-12](#).

TABLE 5-12 SUN-SNMP-NETRA-CT-MIB Trap Forwarding Table

Field	Description
Trap Forwarding Index	A number identifying the Trap forwarding discriminator.
Trap Forwarding Destination	The IP address to which traps identified by this table entry should be sent.
Forwarded Trap ID	The ID of the trap type to which this entry applies. The special value {0 0} indicates that this entry applies to all traps.
Forwarded Trap Object	The object to which this entry applies. By convention, this is the name of the first object in the row in the table referenced. The special value {0 0} indicates that traps of this type from all objects of the type that can generate it. It should also be used when traps from the Netra CT NE are to be specified.
Trap Forwarding Port	The UDP port on the specified management system to which traps identified by this entry should be sent.

TABLE 5-12 SUN-SNMP-NETRA-CT-MIB Trap Forwarding Table

Field	Description
Lowest Forwarded Severity	The lowest severity of traps of this type from the specified object that should be sent to this address. This object has significance only if the trap type specified has a severity associated with it.
Forwarded Indeterminate	When this object has the value <code>TRUE</code> , traps with indeterminate severity will be forwarded to the specified event. This object has significance only if the trap type specified has a severity associated with it.
Trap Forwarding Row Status	This object is used to create a new row or to delete an existing row in the table.

MIB Notification Types

MIB notification types consist of auxiliary definitions for alarms. Except for perceived severity, the objects shown in [TABLE 5-13](#) can be optionally appended to any alarm notification.

TABLE 5-13 MIB Notification Types

Field	Description
Trap Alarm Severity	The perceived severity of the alarm, as specified by the agent that generated it.
Trap Alarm Backed Up	If the value of this object is <code>TRUE</code> , the failed object has been backed up.
Trap Alarm Back-Up Object	Indicates the object that provided back-up services to the failed object.
Trap Alarm Specific Problem	Indicates further refinements to the problem identified by the alarm type. If more than one specific problem is described in this object, the problem descriptions are separated by newline characters.
Trap Alarm Repair Act	Indicates proposed repair actions reported by the agent for the problem identified by the alarm. If more than one action is described in this object, the problem descriptions are separated by newline characters.

SNMP Traps

The SNMP management software has the ability to send traps, or messages, to an application when one or more conditions have been met. Generally, a trap is an unsolicited network packet sent from an agent that usually reports some unexpected error condition.

[TABLE 5-14](#) describes the SNMP traps found in the Netra CT SNMP MIB.

TABLE 5-14 SUN-SNMP-NETRA-CT-MIB Traps

SNMP Trap	Description
netraCtHwHighTempAlarm	Indicates that a high temperature condition has occurred on the hardware unit associated with the specified index.
netraCtHwUnitUp	Indicates that the operational state of the specified hardware unit has transitioned to up.
netraCtHwUnitDown	Indicates that the operational state of the specified hardware unit has transitioned to down.
netraCtHwUnitCreated	Indicates that the specified hardware unit has been installed at the specified location.
netraCtHwUnitDeleted	Indicates that the specified hardware unit has been removed or uninstalled from the specified location.
netraCtInstalledSwCreated	Indicates that the specified software package has been installed.
netraCtInstalledSwDeleted	Indicates that the specified software package has been removed.
netraCtRunningSwCreated	Indicates that the specified software has been started.
netraCtRunningSwDeleted	Indicates that the specified software has been stopped.
netraCtHwMemoryErrorAlarm	Indicates that a memory error has occurred.

TABLE 5-15 defines the standard SNMP traps found in the RFC123-MIB.

TABLE 5-15 RFC1213-MIB Traps

SNMP Trap	Description
coldStart	Signifies that the entity, acting in an agent role, is reinitializing itself and that its configuration might have been altered.
warmStart	Signifies that the entity, acting in an agent role, is reinitializing itself such that its configuration is unaltered.
linkUp	Signifies that the entity, acting in an agent role, has detected that the <code>ifOperStatus</code> object for one of its communication links left the down state and transitioned into some other state (but not into the <code>notPresent</code> state). This other state is indicated by the included value of <code>ifOperStatus</code> .
linkDown	Signifies that the entity, acting in an agent role, has detected that the <code>ifOperStatus</code> object for one of its communication links is about to enter the down state from some other state (but not from the <code>notPresent</code> state). This other state is indicated by the included value of <code>ifOperStatus</code> .

Understanding the MIB Variable Descriptions

TABLE 5-16 defines the MIB elements used in MIB module descriptions in the sections of the MIB file. For detailed information about these elements, refer to the RFC2578 document, which can be downloaded from the <http://www.ietf.org> web site.

Note – Not every MIB element is present for every MIB module.

TABLE 5-16 MIB Variable Syntax

MIB Element	Description
Module name	The name of the MIB module.
Module type	The type of ASN.1 macro used for the module. Macro types are the following: <ul style="list-style-type: none">• OBJECT-TYPE – Defines the type of the managed object.• NOTIFICATION-TYPE – Defines the information contained within an unsolicited transmission of management information (for example, a trap or a request).
SYNTAX	Defines the data structure of the module.

TABLE 5-16 MIB Variable Syntax

MIB Element	Description
MAX-ACCESS	Defines whether the module can read, write, or create an instance of the object, or to include its value in a notification. Can be one of the following: <ul style="list-style-type: none">• <code>not-accessible</code> – Indicates an auxiliary object (objects that are both specified in the INDEX clause of a conceptual row and also columnar objects of the same conceptual row are termed auxiliary objects).• <code>accessible-for-notify</code> – Indicates an object that is accessible only by way of a notification (for example, an SNMP trap).• <code>read-only</code> – Only able to read an instance of the object.• <code>read-write</code> – Able to read and write, but not create an instance of the object.• <code>read-create</code> – Able to read, write, and create an instance of the object provides the maximum level of access (<code>read-create</code> is a superset of <code>read-write</code>).
STATUS	Indicates whether this module definition is current or historic. All of the modules in the SUN-SNMP-NETRA-CT-MIB are current.
DESCRIPTION	Describes the function and use of the module.
INDEX	The INDEX clause defines instance identification information for the columnar objects subordinate to that object. Refer to RFC2578 for more information.
DEFVAL	Defines the default value (DEFVAL) which might be used at the discretion of an SNMP agent when an object instance is created.

For a complete description, see the MIB module in the default location `/opt/SUNWnetract/mgmt2.0/mibs/SUN-SNMP-NETRA-CT-MIB.mib`, delivered as part of the Netra CT software package.

Changing Midplane FRU-ID

This section shows how to change the *locationName* part of FRU-ID.

The Netra CT midplane stores the *locationName*, which is the geographical location of the system, for example, `chassis6`. This value is stored in the alarm card flash and can be set by the customer. The *locationName* enables system monitoring applications to report specific details.

This example uses a NET-SNMP application to interact with MOH and set the midplane's location to a particular value.

1. Determine the index of the midplane object from the `entPhysicalTable`.

At the prompt, type the command:

```
snmpwalk -c public -m SUN-SNMP-NETRA-CT-MIB hostName \
entPhysicalDescr
```

where:

-c *community* specifies the community string.

-m SUN-SNMP-NETRA-CT-MIB specifies that the Netra CT MIB should be loaded.

hostName is the development system running MOH.

This process and its results are shown in [CODE EXAMPLE 5-1](#)

CODE EXAMPLE 5-1 Index of the Midplane Object

```
$snmpwalk -c public -m SUN-SNMP-NETRA-CT-MIB hostName:9161 entPhysicalDescr
ENTITY-MIB::entPhysicalDescr.2 = STRING: 01ae 5405026 Midplane 0000
ENTITY-MIB::entPhysicalDescr.3 = STRING: scb_slot
ENTITY-MIB::entPhysicalDescr.4 = STRING: fan_slot
ENTITY-MIB::entPhysicalDescr.5 = STRING: fan_slot
ENTITY-MIB::entPhysicalDescr.6 = STRING: ps_slot
ENTITY-MIB::entPhysicalDescr.7 = STRING: crtm_slot
ENTITY-MIB::entPhysicalDescr.8 = STRING: cftm_slot
ENTITY-MIB::entPhysicalDescr.9 = STRING: cpci_slot
ENTITY-MIB::entPhysicalDescr.10 = STRING: cpci_slot
ENTITY-MIB::entPhysicalDescr.11 = STRING: cpci_slot
ENTITY-MIB::entPhysicalDescr.12 = STRING: cpci_slot
ENTITY-MIB::entPhysicalDescr.13 = STRING: cpci_slot
ENTITY-MIB::entPhysicalDescr.14 = STRING: prtm_slot
ENTITY-MIB::entPhysicalDescr.15 = STRING: pdu
ENTITY-MIB::entPhysicalDescr.25 = STRING: 01ae 5016118 scb 0499
ENTITY-MIB::entPhysicalDescr.26 = STRING: ssp_slot
ENTITY-MIB::entPhysicalDescr.27 = STRING: ssp
ENTITY-MIB::entPhysicalDescr.28 = STRING: 01ae 5404931 fan 0499
ENTITY-MIB::entPhysicalDescr.29 = STRING: 01ae 5404931 fan 0499
ENTITY-MIB::entPhysicalDescr.30 = STRING: 01ae 3001535 ps 0399
ENTITY-MIB::entPhysicalDescr.31 = STRING: cftm
ENTITY-MIB::entPhysicalDescr.32 = STRING: 0000 5016123 0101 0000
ENTITY-MIB::entPhysicalDescr.33 = STRING: RJ45
ENTITY-MIB::entPhysicalDescr.35 = STRING: RJ45
ENTITY-MIB::entPhysicalDescr.37 = STRING: RJ45
ENTITY-MIB::entPhysicalDescr.39 = STRING: RJ45
ENTITY-MIB::entPhysicalDescr.41 = STRING: DB15
```

2. Set the midplane location to the new value of `chassis6` using the following command:

```
$snmpset -c public -m SUN-SNMP-NETRA-CT-MIB hostName:9161 \  
netraCtEquipLocation.1 = chassis6
```

3. Show the current value of the midplane's location.

At the prompt, type the command:

```
$snmpget -c public -m SUN-SNMP-NETRA-CT-MIB hostName:9161 \  
netraCtEquipLocation.1
```

The result displays the identifying string of the location of any Netra CT equipment locations, as shown in [CODE EXAMPLE 5-2](#).

CODE EXAMPLE 5-2 Identifying the Midplane's Current Location

```
$snmpget -c public -m SUN-SNMP-NETRA-CT-MIB hostName:9161 \  
netraCtEquipLocation.1  
SUN-SNMP-NETRA-CT-MIB::netraCtEquipLocation.1 = STRING: chassis6
```

Setting High Temperature Alarms

An alarm in SNMP is defined as a trap with a severity associated with it. When a `HIGH_TEMPERATURE` alarm (CPU high temperature) occurs, the user's application will receive the SNMP trap `netraCtHwHighTempAlarm`, and `netraCtIfChanged` trap for the `ifOperStatus` of the interface corresponding to the alarm output port. The user's application also will receive alarm clear traps when the condition of alarms are cleared, and an attribute change trap of the `ifOperStatus`.

The Netra CT alarm card supports three output alarm interfaces. The alarm pins (`alarm0`, `alarm1`, `alarm2`) are statically mapped into severities of critical, major, minor respectively. When an alarm occurs, the corresponding alarm pin is driven high according to the severity of the alarm.

The following example shows how to set the high temperature alarm from the default to major.

▼ To Set the High Temperature Alarm Severity to Major

1. Create an entry in the `netraCtAlarmSevProfileTable`.

At the prompt, type the command:

```
$snmpset -c public -m SUN-SNMP-NETRA-CT-MIB hostName\  
netraCtAlarmSevProfileRowStatus.1 = 4
```

where:

`-c community` specifies the community string.

`-m SUN-SNMP-NETRA-CT-MIB` specifies that the Netra CT MIB should be loaded.

`hostName` is the development system running MOH.

This process and its result are shown in [CODE EXAMPLE 5-3](#).

CODE EXAMPLE 5-3 Creating an Entry in the Profile Table

```
$snmpset -c public -m SUN-SNMP-NETRA-CT-MIB localhost:9161 \  
netraCtAlarmSevProfileRowStatus.1 = 4  
SUN-SNMP-NETRA-CT-MIB::netraCtAlarmSevProfileRowStatus.1 = INTEGER: active(1)
```

Creating an entry in the `netraCtAlarmSevProfileTable` also creates an entry in the `netraCtAlarmSevTable`. The entry in the latter corresponds to the profile entry and translates the high temperature alarm entry into the row of integers shown in [CODE EXAMPLE 5-4](#).

CODE EXAMPLE 5-4 Automatic Entry Created in Corresponding Alarm Severity Table

```
$snmpwalk -c public -m SUN-SNMP-NETRA-CT-MIB localhost:9161\  
netraCtAlarmSevTable  
SUN-SNMP-NETRA-CT-MIB:\br/>:netraCtAlarmSeverity.1.15.1.3.6.1.4.1.42.2.65.1.1.1.2.0.34 = INTEGER:\br/>minor(3)  
End of MIB
```

2. Set the severity of `netraCtHighTempAlarm` for this profile.

At the prompt, type the command:

```
$ snmpset -c public -m SUN-SNMP-NETRA-CT-MIB localhost:9161\  
netraCtAlarmSeverity.1.15.1.3.6.1.4.1.42.2.65.1.1.1.2.0.34 = 2
```

where:

1.3.6.1.4.1.42.2.65.1.1.1.2.0.34 represents the string
'netraCtHighTempAlarm'

The entry at = (in this example, 2) establishes a major alarm severity.

The result is shown in [CODE EXAMPLE 5-5](#).

CODE EXAMPLE 5-5 Setting the Alarm Severity for the Profile Table

```
$ snmpset -c public -m SUN-SNMP-NETRA-CT-MIB localhost:9161\  
netraCtAlarmSeverity.1.15.1.3.6.1.4.1.42.2.65.1.1.1.2.0.34 = 2  
SUN-SNMP-NETRA-CT-MIB\  
:netraCtAlarmSeverity.1.15.1.3.6.1.4.1.42.2.65.1.1.1.2.0.34 = INTEGER:  
major(2)
```

3. Set `netraCtEquipAlarmSeverityIndex` of the thermistor entry to correspond with the `netraCtAlarmSevProfile` entry from the `netraCtAlarmSevProfileTable`.

At the prompt, type the command:

```
$ snmpset -c public -m SUN-SNMP-NETRA-CT-MIB localhost:9161 \  
netraCtEquipAlarmSeverityIndex.2 = 1
```

This example uses the `netraCtAlarmSevProfileTable` entry from [CODE EXAMPLE 5-3](#). The index of that entry was the integer 1 in the statement: `netraCtAlarmSevProfileRowStatus.1`. The result of this process is shown in [CODE EXAMPLE 5-6](#).

CODE EXAMPLE 5-6 Setting the Index Entry Corresponding to the Thermistor

```
$snmpset -c public -m SUN-SNMP-NETRA-CT-MIB localhost:9161 \  
netraCtEquipAlarmSeverityIndex.2 = 1  
SUN-SNMP-NETRA-CT-MIB::netraCtEquipAlarmSeverityIndex.2 = INTEGER: 1
```

When the CPU temperature returns to normal, the alarms are cleared automatically. For further information, refer to [“SUN-SNMP-NETRA-CT-MIB”](#) on page 55.

Managed Object Hierarchy Software Modules

This chapter provides a high-level description of the Release 2 Management Object Hierarchy (MOH) software modules for the Netra CT platform. It describes the software module interfaces and their major internal modules. It consists of:

- [“Software Module Design” on page 74](#)
- [“Software Services” on page 74](#)
- [“Software Module MBeans” on page 75](#)

Software Module Design

The software services in the system are monitored by software modules that are part of the Information Module layer objects.

Software services are either reliable services (such as RNFS, RBS, or CGTP) or unreliable services (such as TFTP, or NIS). The software services can be a software subsystem such as a network stack (TCP, IP, UDP); an I/O driver such as a network driver; or network processes or network daemons such as NFS.

Some software services are only available on certain CPU boards. For example, CGTP is available for both the host and the satellite CPU boards, but RBS or RNFS are only available on the host CPU board.

The software module interacts with the OS platform through Java interfaces to:

- Monitor OS platform software services for software status, such as installed or not installed and configured or not configured
- Monitor software service subsystems and daemons for status, such as running or not running
- Provide traps and notifications for events related to the status of software services

The software module also provides APIs for management applications to configure the monitoring of software services, such as setting error thresholds, setting polling intervals, starting and stopping polling, and setting maximum retry-counts for the recovery of the daemons.

Software Services

The software modules monitor the following software services:

TABLE 6-1 Software Services

Software Module	Software Service Monitored
CGTP	RDHCP
Ethernet Interface Statistics	RNFS
NIS	SNDR
PMS	TCP, IP, UDP

TABLE 6-1 Software Services

Software Module	Software Service Monitored
Platform Management Service (PICLD on Solaris OS)	TFTP
RDHCP	UFS
RNFS	

Software Module MBeans

This section describes the software module MBeans for each of the software services that the MOH software modules monitor. As specified by the Java Management Extensions for instrumentation, all attributes and operations are explicitly listed in an MBean interface. This interface must have the same name as the class that implements it, followed by the MBean suffix. Since the interface and its implementation are usually in different files, there are two files which make up a standard MBean.

These MBeans and their public APIs provide the management interface to manage the applications. All the specific MBeans below are extended from the `SoftwareServiceMBean`. For more specific information, refer to the Java documents for the APIs that are distributed as part of the Netra CT MOH package. See [“Viewing the Netra CT Management Agent API Online”](#) on page 44 for details.

SoftwareMonitorMBean

The `SoftwareMonitorMBean` is an object that clients can use to discover all the software services in the system. The `SoftwareMonitorMBean` contains the method `getSoftwareServiceList()` which returns the list of software services.

DaemonMBean

This class provides the name of the daemon, the state of the daemon, and the daemon recovery try count.

SoftwareServiceMBean

The `SoftwareServiceMBean` provides the base class from which other `ServiceMBeans` are extended. The `SoftwareServiceMBean` provides the following:

- Name of the services
- Status of the service (up or down)
- Getting and setting polling intervals
- Starting the polling
- Stopping the polling
- Getting and setting the number of excessive error intervals. This number is the threshold that determines if an event is sent to a client. If an error count exceeds this number, an error event is sent. There will be no more error events until the error condition disappears or a clear event is sent. For example, assume that the error threshold is set at 5% error per total transaction and the number of excessive intervals is set at 3. If the error exceeds 5% in more than three consecutive polling intervals, a file system error event is sent to the client.
- Getting a list of `DaemonMBeans` that support the service, if any

NfsServiceMBean

The `NfsServiceMBean` enables the client to monitor the NFS services. A client can get and set the maximum error threshold, get and set the threshold for excessive error intervals, and get the list of NFS mount failures.

UfsServiceMBean

The `UfsServiceMBean` enables the client to monitor the UFS services. A client can get and set the maximum threshold of the file system usage percentage, get and set the threshold for the number of excessive usage intervals, and query the list of file systems exceeding the usage threshold.

TcpServiceMBean

The `TcpServiceMBean` enables the client to monitor the TCP services. A client can get status and statistics for the TCP network layer, get and set intervals and thresholds for gathering the statistics, start and stop polling, and get a list of daemons supporting the service.

UdpServiceMBean

The `UdpServiceMBean` enables the client to monitor the UDP services. A client can get status and statistics for the UDP network layer, get and set intervals and threshold for gathering the statistics, start and stop polling, and get a list of daemons supporting the service.

IpServiceMBean

The `IpServiceMBean` enables the client to monitor the IP services. A client can get status and statistics for the IP network layer, get and set intervals and thresholds for gathering the statistics, start and stop polling, and get a list of daemons supporting the service.

EtherIfStatsMBean

The `EtherIfStatsMBean` monitors the Ethernet drivers, and monitors the interface for transmitter and receiver error counts. A client can set and get the maximum error threshold, set and get the threshold for number of excessive intervals, and query for the list of Ethernet interfaces in error.

CgtpServiceMBean

The `CgtpServiceMBean` enables the client to monitor the CGTP services. A client can get status and statistics for the IP network layer, list and get state of associated Ethernet physical interfaces, get and set intervals and thresholds for gathering the statistics, start and stop polling, and get a list of daemons supporting the service.

RnfsServiceMBean

The `RnfsServiceMBean` enables the client to monitor the RNFS services. A client can get status and statistics for the UDP network layer, get and set intervals and thresholds for gathering the statistics, start and stop polling, and get a list of daemons supporting the service.

Processor Management Services

This chapter describes the processor management services (PMS) application programming interface (API). This chapter contains the following sections:

- [“PMS Software Overview” on page 80](#)
- [“PMS Man Pages” on page 83](#)
- [“PMS Examples” on page 84](#)

PMS Software Overview

The Processor Management Services (PMS) software is an extension to the Netra CT platform services software that addresses the requirements of high-availability (HA) application frameworks. The PMS software enables client applications to manage the operation of the processor nodes within a single Netra CT system or within a cluster of multiple Netra CT systems. A processor node is a combination of CPU blade hardware, CPU memory, I/O interfaces, the operating system that runs on them, and select applications. A PMS cluster can include the alarm card and all of the CPU cards in a single Netra CT system, or it can include a defined group of alarm cards and CPU cards located in multiple systems.

The PMS software provides distributed CPU board resource management infrastructure for clusters of CPU boards. This infrastructure includes low-level administrative control and monitoring, high-level configuration, fault recovery, and user-interface functionality. [FIGURE 7-1](#) identifies the architectural components of the Netra CT software services.

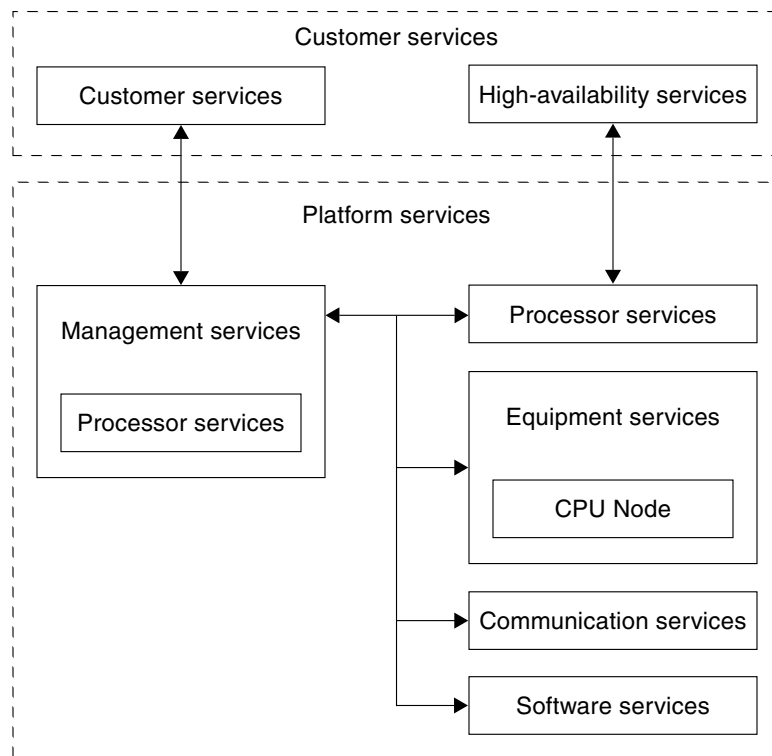


FIGURE 7-1 Netra CT Software Services

In a Netra CT cluster, the PMS software runs on both the alarm cards and the CPU boards. The PMS software running on alarm cards provides local and remote service connections for managing the CPU cards in its system. The PMS software running on CPU cards provides local and remote service connections for managing the resources running on the board, and the software provides remote access for managing resources running on other CPU cards in a PMS cluster.

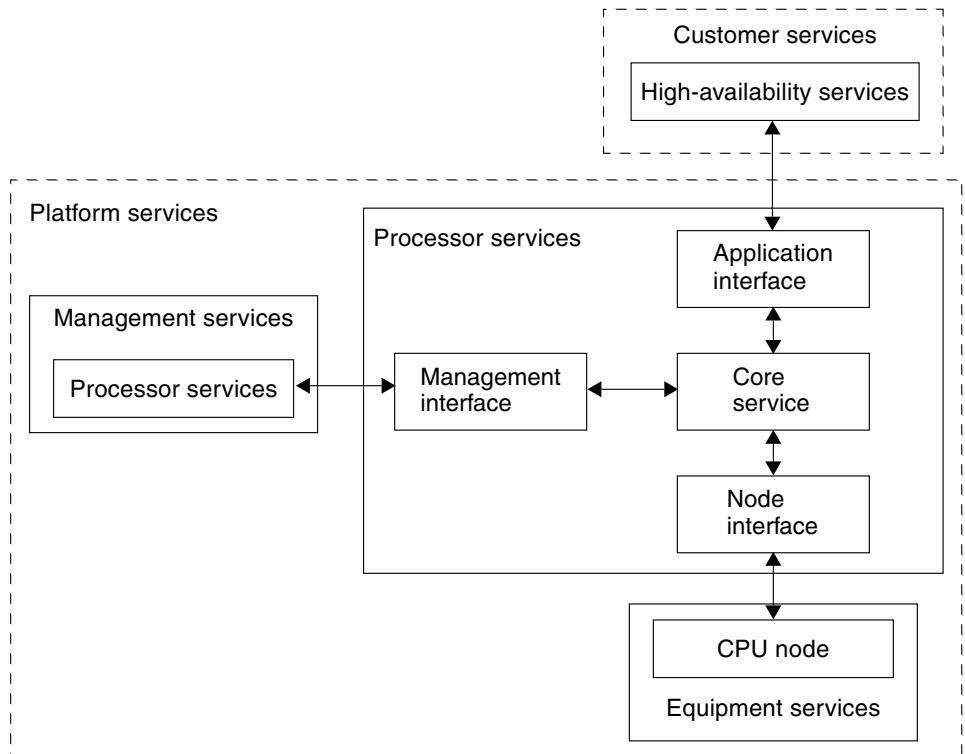


FIGURE 7-2 PMS Software Services and Interfaces

[FIGURE 7-2](#) indicates the internal interfaces of the processor services.

The PMS software organizes the CPU resources it manages into the following three groups:

- Resource group 0 (RG0) – Specific application services
- Resource group 1 (RG1) – Operating system functionality
- Resource group 2 (RG2) – CPU hardware and the remaining processor board resources

The PMS software that runs on both alarm cards and the CPU cards divides its functionality along client-side and server-side (daemon-side) lines. The common client-side function provides a shared API for up to eight simultaneous application service processes. The core API functionality includes API control, PMS daemon control, application PMS connectivity, and application message send and receive with function execution. The API provides per-process serialization and separate threads for message reception and user-defined function execution, and messaging process timing.

In a typical example, a PMS client detects resource failures remotely and then remotely activates replacement resources such as those found in high-availability applications. The common daemon function provides server-side control and monitoring functionality for up to 16 remote CPUs. The daemon function also provides client-side functionality for controlling and monitoring up to 16 remote CPUs simultaneously with minimized latency by way of per-remote-CPU threading, as well as daemon control and performance monitoring and resource group monitoring and control.

From the client side, the alarm card function available by way of the send and receive messaging API is broken into management and drawer blocks. (The PMS software refers to Netra CT systems as drawers.) The CPU cards are divided into management node and remote node drawer (RND) views. The management view on both the alarm card and the CPU board provides administrative control and status over the PMS daemon as a whole. The management view also monitors the PMS software's performance.

The drawer (system) view by means of the alarm card provides the following administrative controls and monitors of the RG2 (hardware) resources: Core power down, power up, and reset. For RG1 (operating system) resources, this view also provides the following administrative controls and monitors: core shutdown, boot, and reboot. For RG0 (application services), this view provides off-line and active administrative controls. Finally, for the combined resource groups, this view provides the following administrative controls and monitors: Core maintenance, and operational configuration, five recovery processes, and the graceful reboot of the group.

The node view, by way of the CPU card itself, provides a much reduced set of administrative controls and monitors relative to the drawer view of the hardware, operating system, and the same administrative controls and monitors of the application services. In RG2 only reset administrative controls exist, but no monitors. Likewise, in RG1 only reboot administrative controls exist, but no monitors. In this view, there is no administrative control over the combined resource groups.

The CPU card RND view provides remote system view administrative controls and monitors to all the resource groups, with the exception of an alarm card failure. In this failure case, a reduced remote node view is used.

The PMS software execution performance is targeted by scheduling optimizations as well as using lightweight, proprietary messaging protocols, intersystem data encoding, and packetization protocols. The PMS software scalability due to CPU card growth is addressed by a per-CPU multithreading of up to 16 remote CPU cards per CPU. Application client growth is addressed by way of per-process multithreading with up to eight client processes per PMS daemon.

The PMS software performance and reliability in cluster communication is also addressed with a messaging infrastructure that supports unidirectional and bidirectional point-to-point and unidirectional point-multipoint channels. This infrastructure includes source time-stamping available to the client for latency detection, call and return time-out for failure detection, and interprocess and intersystem TCP/IP socket streams for connection control, reachability determination, and reliable transport.

PMS Man Pages

The PMS software application programming interface (API) has been documented completely in the UNIX man pages included with the Netra CT software. [TABLE 7-1](#) lists the man pages included with the Netra CT PMS software:

TABLE 7-1 Processor Management Services Man Pages

Man page	Description
<code>pms(1M)</code>	Provides an overview of the PMS software.
<code>pmsd(1M)</code>	Describes how to start and stop the CPU board PMS daemon (<code>pmsd</code>) and lists the daemon's command line options.
<code>pmsd_ac(1M)</code>	Describes how to start and stop the alarm card PMS daemon (<code>pmsd_ac</code>) from the command line interface, and lists all the daemon's other command-line functions.
<code>pms_apistart(1M)</code>	Describes the PMS API functions used to initialize (<code>pms_apistart</code>) and to free up (<code>pms_apistop</code>) PMS API resources in a PMS process. The man page also documents the functions used to take PMS out of an inactive state (<code>pms_start</code>) and to return it to an inactive state (<code>pms_stop</code>).

TABLE 7-1 Processor Management Services Man Pages (*Continued*)

Man page	Description
<code>pms_connect(1M)</code>	Documents the PMS API functions used to create (<code>pms_connect</code>) and destroy (<code>pms_disconnect</code>) a PMS daemon interface session.
<code>pms_send(1M)</code>	Describes the PMS API functions that enable PMS clients to send (<code>pms_send</code>) and receive (<code>pms_receive</code>) messages with other PMS clients or clusters.
<code>pms_usermgmt_message_payloads(1M)</code>	Describes the payloads for the user and management PMS function groups.
<code>pms_node_message_payloads(1M)</code>	Defines the payloads for the node PMS function group.
<code>pms_rnd_message_payloads(1M)</code>	Describes the payloads for the remote node drawer (system) PMS function group.

If you cannot view these man pages, add the PMS man page directory location to your `$MANPATH` environment variable. By default, the PMS man pages are installed in the following directory: `/opt/SUNWnetract/mgmt2.0/man`. Depending on the UNIX shell you are using, this variable might be defined in a shell startup file. Refer to the Solaris documentation for instructions on adding the PMS man page directory to a UNIX shell startup file on your system.

PMS Examples

The following examples show how to initialize a PMS client, the structure of the main thread, asynchronous messaging, scheduling, and the PMS client's user and management, node, and RND interfaces.

- [“PMS Client Initialization Example” on page 85](#)
- [“PMS Client Main Thread” on page 91](#)
- [“PMS Client Asynchronous Message Handling” on page 93](#)
- [“PMS Client Scheduling Example” on page 106](#)
- [“PMS Client User and Management Interface” on page 107](#)
- [“PMS Client Node Interface” on page 124](#)
- [“PMS Client RND Interface” on page 131](#)

CODE EXAMPLE 7-1 begins by initializing the main thread for a PMS client.

CODE EXAMPLE 7-1 PMS Client Initialization Example

```
#include <sys/types.h>          /* socketpair() */
#include <sys/socket.h>         /* socketpair() */

#include <unistd.h>             /* write(), read() */

#include <signal.h>            /* sigemptyset(), sigaddset(), sigaction() */
#include <time.h>              /* timer_create(), timer_settime() */

#include <stdio.h>             /* printf(), scanf() */

#include "pms.h"

/* Application State Machine Example Overview:
   1) PMS API initialization and usage.
   2) PMS Daemon connectivity and availability management.
   3) Named application synchronization and behavior.
   4) Remote Node Drawer address list synchronization and monitoring.
   5) Basic example data caching synchronization on the client side for PMS
      items a particular application's intent/design makes it interested in.
   6) Basic asynchronous message handling infrastructure for the application.
   7) Remote monitoring of remote node drawer's.
   8) Example Control of a pair of remote node drawer's(not implemented yet).
*/

void* app_hasim_thread(void*);

/* Event message handlers.. */

/* This mechanism registers one receive handler with PMS for all messages, which
   simply posts the messages to the client thread's processing queue to have them
   handled synchronously. Alternatively, handlers can be registered with PMS
   individually in which case they will execute asynchronous to the client thread
   in the context of the PMS API receive thread. */

void app_hasim_receive_post(struct pms_receive *pr);
int app_hasim_receive_dispatch(struct pms_receive* pr);

void app_hasim_receive_user_status(struct pms_receive *pr);
void app_hasim_receive_mgmt_status(struct pms_receive *pr);
```

CODE EXAMPLE 7-1 PMS Client Initialization Example (Continued)

```
void    app_hasim_receive_node_rg0_status(struct pms_receive *pr);
void    app_hasim_receive_node_rg0_app_state_set_execute\
        (struct pms_receive *pr);
void    app_hasim_receive_rnd_status(struct pms_receive *pr);
void    app_hasim_receive_rnd_md0_status(struct pms_receive *pr);

void    app_hasim_receive_time_status(void);

/* Convenient state machine process sub-groupings.. */

void    app_hasim_user_process(void);
void    app_hasim_mgmt_process(void);
void    app_hasim_node_process(void);
void    app_hasim_rnd_process(void);
void    app_hasim_process(void);

/* Timer signal handler.. */

void    app_hasim_sigusr1_signal_handler(int);

/* Interval's currently set for example convenience.. */

#define HASIM_CHECK_INTERVAL                2
#define HASIM_SYNCHECK_INTERVAL            600
#define HASIM_CHECK_VALID_INTERVAL         1800
#define HASIM_CHECK_INVALID_INTERVAL       3600

#define HASIM_RND_ADDRESS_AUDIT_ENTRIES    2

struct hasim_info
{
    int                sockfd[2];
    struct
    {
        char          node_ip_address[20];
        char          drawer_ip_address[20];
        int           node_slot_number;
    } rnd_address[HASIM_RND_ADDRESS_AUDIT_ENTRIES];

    struct
    {
#define HASIM_USER_RECEIVE_UNREGISTERED    0x00
#define HASIM_USER_RECEIVE_REGISTERED     0x01
        int           receive_state;
#define HASIM_USER_PMS_VIEW_REACHABLE     0x00
```

CODE EXAMPLE 7-1 PMS Client Initialization Example (Continued)

```

#define HASIM_USER_PMS_VIEW_UNREACHABLE 0x01
    int pms_view;

    int view_cache;
} user_info;

    struct
    {
#define HASIM_MGMT_RECEIVE_UNREGISTERED 0x00
#define HASIM_MGMT_RECEIVE_REGISTERED 0x01
        int receive_state;
#define HASIM_MGMT_PMS_STATE_UNAVAILABLE 0x00
#define HASIM_MGMT_PMS_STATE_AVAILABLE 0x01
        int pms_state;
#define HASIM_MGMT_RND_ADDRESS_UNVERIFIED 0x00
#define HASIM_MGMT_RND_ADDRESS_VERIFIED 0x01
        int rnd_address_state;
        int rnd_address_identifier[16];

#define HASIM_MGMT_CACHE_INVALID 0x00
#define HASIM_MGMT_CACHE_OLD 0x01
#define HASIM_MGMT_CACHE_VALID 0x02
        int cache_state;
        int last_update;
        int last_sync_check;

        int mgmt_state_cache;
        struct
        {
            int identifier;
            char node_ip_address[20];
            char drawer_ip_address[20];
            int node_slot_number;
        } rnd_address_cache[16];
    } mgmt_info;

    struct
    {
#define HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED 0x02
#define HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED 0x04
#define HASIM_NODE_GROUP_RECEIVE_UNREGISTERED 0x00
#define HASIM_NODE_GROUP_RECEIVE_REGISTERED 0x06
        int receive_state;
#define HASIM_NODE_RG0_APP_NAME_UNREGISTERED 0x00
#define HASIM_NODE_RG0_APP_NAME_REGISTERED 0x01
        int rg0_app_name_state;
#define HASIM_NODE_SERVICE_STATE_OFFLINE 0x00

```

CODE EXAMPLE 7-1 PMS Client Initialization Example (Continued)

```
#define HASIM_NODE_SERVICE_STATE_ACTIVE 0x01
    int service_state;

#define HASIM_NODE_CACHE_INVALID 0x00
#define HASIM_NODE_CACHE_OLD 0x01
#define HASIM_NODE_CACHE_VALID 0x02
    int cache_state;
    int last_update;
    int last_sync_check;

    int rg0_state_cache;
} node_info;

struct
{
#define HASIM_RND_RECEIVE_REGISTERED 0x01
#define HASIM_RND_MD0_RECEIVE_REGISTERED 0x20
#define HASIM_RND_GROUP_RECEIVE_UNREGISTERED 0x00
#define HASIM_RND_GROUP_RECEIVE_REGISTERED 0x21
    int receive_state;

#define HASIM_RND_CACHE_INVALID 0x00
#define HASIM_RND_CACHE_OLD 0x01
#define HASIM_RND_CACHE_VALID 0x02
    int cache_state;
    int last_update;
    int last_sync_check;

    int view_cache;
    int md0_config_cache;
} rnd_info[16];

};

static struct hasim_info mdi;

int
main(int argc, char *argv[])
{
    struct pms_receive pr;
    struct sigaction sigusr1_signal_handler_info;
    struct sigevent evp;
    timer_t timerid;
    struct itimerspec val;
    struct itimerspec oval;
```


CODE EXAMPLE 7-1 PMS Client Initialization Example (*Continued*)

```
int                i;

if (argc != 1)
{
    printf("Invalid Arguments\n");

    exit(1);
}

/* Start/Initialize the PMS API before using any further calls.. */

if (pms_apistart() == -1)
    exit(2);

/* Create message queue.. */

if (socketpair(AF_UNIX, SOCK_DGRAM, 0, mdi.sockfd) == -1)
{
    exit(3);
}

/* Setup defaults.. */

/* Audit DB hardcoding for this example.. */

strcpy(&mdi.rnd_address[0].node_ip_address[0], "129.150.94.70");
strcpy(&mdi.rnd_address[0].drawer_ip_address[0], "129.150.151.140");
mdi.rnd_address[0].node_slot_number = 2;
strcpy(&mdi.rnd_address[1].node_ip_address[0], "129.150.94.58");
strcpy(&mdi.rnd_address[1].drawer_ip_address[0], "129.150.151.143");
mdi.rnd_address[1].node_slot_number = 3;

mdi.user_info.receive_state = HASIM_USER_RECEIVE_UNREGISTERED;
mdi.user_info.pms_view = HASIM_USER_PMS_VIEW_UNREACHABLE;

mdi.mgmt_info.receive_state = HASIM_MGMT_RECEIVE_UNREGISTERED;
```

CODE EXAMPLE 7-1 PMS Client Initialization Example (Continued)

```
mdi.mgmt_info.pms_state = HASIM_MGMT_PMS_STATE_UNAVAILABLE;
mdi.mgmt_info.rnd_address_state = HASIM_MGMT_RND_ADDRESS_UNVERIFIED;
for(i=0;i<16;i++)
    mdi.mgmt_info.rnd_address_identifer[i] = -1;
mdi.mgmt_info.cache_state = HASIM_MGMT_CACHE_INVALID;
mdi.mgmt_info.last_update = HASIM_CHECK_INVALID_INTERVAL;
mdi.mgmt_info.last_sync_check = HASIM_SYNCCHECK_INTERVAL;

mdi.node_info.receive_state = HASIM_NODE_GROUP_RECEIVE_UNREGISTERED;
mdi.node_info.rg0_app_name_state = HASIM_NODE_RG0_APP_NAME_UNREGISTERED;
mdi.node_info.service_state = HASIM_NODE_SERVICE_STATE_OFFLINE;
mdi.node_info.cache_state = HASIM_NODE_CACHE_INVALID;
mdi.node_info.last_update = HASIM_CHECK_INVALID_INTERVAL;
mdi.node_info.last_sync_check = HASIM_SYNCCHECK_INTERVAL;

for(i=0;i<16;i++)
{
    mdi.rnd_info[i].receive_state = HASIM_RND_GROUP_RECEIVE_UNREGISTERED;
    mdi.rnd_info[i].cache_state = HASIM_RND_CACHE_INVALID;
    mdi.rnd_info[i].last_update = HASIM_CHECK_INVALID_INTERVAL;
    mdi.rnd_info[i].last_sync_check = HASIM_SYNCCHECK_INTERVAL;
}

/* Setup timer.. */

sigemptyset(&sigusr1_signal_handler_info.sa_mask);
sigaddset(&sigusr1_signal_handler_info.sa_mask, SIGUSR1);
sigusr1_signal_handler_info.sa_flags = 0;
sigusr1_signal_handler_info.sa_handler = app_hasim_sigusr1_signal_handler;
sigaction(SIGUSR1, &sigusr1_signal_handler_info, NULL);

evp.sigev_notify = SIGEV_SIGNAL;
evp.sigev_signo = SIGUSR1;

if (timer_create(CLOCK_REALTIME, &evp, &timerid) == -1)
    exit(4);

val.it_value.tv_sec = HASIM_CHECK_INTERVAL;
val.it_value.tv_nsec = 0;
val.it_interval.tv_sec = HASIM_CHECK_INTERVAL;
val.it_interval.tv_nsec = 0;

if (timer_settime(timerid, TIMER_RELTIME, &val, NULL) == -1)
    exit(4);
```

CODE EXAMPLE 7-1 PMS Client Initialization Example (*Continued*)

```
/* Don't bother creating another thread, run in context of main default.. */
app_hasim_thread(0);

}
```

CODE EXAMPLE 7-2 PMS Client Main Thread

```
void*
app_hasim_thread(void* arg)
{
    char                receivebuffer[256];
    int                 receivestatus;

    fd_set              readfds;
    int                 select_return;
    struct timeval      timeout;

    struct pms_send     ps;
    struct pms_receive  pr;

    int                 i;

    printf("*** HA Client Application Simulation ***\n");

    /* Presuming PMS will have been started at boot or by another app.. */

    timeout.tv_sec = HASIM_CHECK_INTERVAL;
    timeout.tv_usec = 0;

    while(1)
    {
        FD_ZERO(&readfds);

        FD_SET(mdi.sockfd[1], &readfds);

        /* Wait for event messages.. */
```

CODE EXAMPLE 7-2 PMS Client Main Thread (*Continued*)

```
select_return = select(64, &readfds, NULL, NULL, &timeout);

if (select_return > 0)
{
    if (FD_ISSET(mdi.sockfd[1], &readfds) != 0)
    {
        receivestatus = read(mdi.sockfd[1], &receivebuffer[0], 256);

        if (receivestatus <= 0)
        {
            /* Handle Error.. */

        }
        else
        {
            /* Handle Message.. */

            app_hasim_receive_dispatch((struct pms_receive*)&receivebuffer[0]);

        }
    }
}
else if (select_return == 0)
{
    /* Handle Timeout.. */

}
else
{
    /* Handle Error.. */

}
}

void
app_hasim_sigusr1_signal_handler(int signal)
{
```

CODE EXAMPLE 7-2 PMS Client Main Thread (*Continued*)

```
struct pms_receive pr;

pr.session.type = PMS_SR_CALL_NO_RETURN;
pr.payload.type = PMS_PD_PAYLOAD_TYPE_MAX+1;

app_hasim_receive_post(&pr);
}
```

The following example sets up a PMS client to handle asynchronous messages.

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling

```
void
app_hasim_receive_post(struct pms_receive* pr)
{

    int                status;

    /* Write for reading in context of main thread.. */
    status = write(mdi.sockfd[0], pr, sizeof(struct pms_receive));

    if (status < 0)
        {
        }

}

int
app_hasim_receive_dispatch(struct pms_receive* pr)
{

    switch(pr->payload.type)
        {
        case PMS_PD_USER_STATUS:

            app_hasim_receive_user_status(pr);

        }

}
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
break;

case PMS_PD_MGMT_STATUS:

    app_hasim_receive_mgmt_status(pr);

break;

case PMS_PD_NODE_RG0_STATUS:

    app_hasim_receive_node_rg0_status(pr);

break;

case PMS_PD_NODE_RG0_APP_STATE_SET_EXECUTE:

    app_hasim_receive_node_rg0_app_state_set_execute(pr);

break;
case PMS_PD_RND_STATUS:

    app_hasim_receive_rnd_status(pr);

break;
case PMS_PD_RND_MD0_STATUS:

    app_hasim_receive_rnd_md0_status(pr);

break;
case PMS_PD_PAYLOAD_TYPE_MAX+1:

    app_hasim_receive_time_status();

break;
}

return(0);

}

void
app_hasim_receive_user_status(struct pms_receive* pr)
{
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
switch(pr->payload.data.user_status.code)
{
case PMS_PD_USER_STATUS_PMS_REACHABLE:

    printf("hasim :          received USER_STATUS PMS_REACHABLE..\n");

    mdi.user_info.view_cache = PMS_PD_USER_STATUS_PMS_REACHABLE;

    /* Run state machine.. */

    app_hasim_process();

break;
case PMS_PD_USER_STATUS_PMS_UNREACHABLE:

    printf("hasim :          received USER_STATUS PMS_UNREACHABLE..\n");

    mdi.user_info.view_cache = PMS_PD_USER_STATUS_PMS_UNREACHABLE;

    app_hasim_process();

break;
}
}

void
app_hasim_receive_mgmt_status(struct pms_receive* pr)
{
    struct pms_send      ps;
    struct pms_receive   prs;

    int                  info_get_fail;

    int                  rnd_address_identifier[16];
    char                 rnd_address_node_ip_address[16][20];
    char                 rnd_address_drawer_ip_address[16][20];
    int                  rnd_address_node_slot_number[16];
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
int                i, j;

switch(pr->payload.data.mgmt_status.code)
{
  case PMS_PD_MGMT_STATUS_PMS_STATE_AVAILABLE:

    printf("hasim :      received MGMT_STATUS PMS STATE AVAILABLE..\n");

    /* Update cached data and set update time.. */

    mdi.mgmt_info.mgmt_state_cache = PMS_PD_MGMT_INFO_GET_STATUS_AVAILABLE;
    mdi.mgmt_info.last_update = 0;

    app_hasim_process();

  break;
  case PMS_PD_MGMT_STATUS_PMS_STATE_UNAVAILABLE:

    printf("hasim :      received MGMT_STATUS PMS STATE UNAVAILABLE..\n");

    mdi.mgmt_info.mgmt_state_cache = PMS_PD_MGMT_INFO_GET_STATUS_UNAVAILABLE;
    mdi.mgmt_info.last_update = 0;

    app_hasim_process();

  break;
  case PMS_PD_MGMT_STATUS_PMS_ADMIN_STATE_FORCE_UNAVAILABLE:

    printf("hasim :      received MGMT_STATUS PMS ADMIN STATE FORCE\
          UNAVAILABLE..\n");

    /* Doing nothing at the moment.. */

  break;
  case PMS_PD_MGMT_STATUS_PMS_ADMIN_STATE_VOTE_AVAILABLE:

    printf("hasim :      received MGMT_STATUS PMS ADMIN STATE VOTE AVAILABLE\
          ..\n");

    /* Doing nothing at the moment.. */

  break;
  case PMS_PD_MGMT_STATUS_PMS_ADMIN_STATE_FORCE_AVAILABLE:
```


CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
printf("hasim :          received MGMT_STATUS PMS ADMIN STATE FORCE \
AVAILABLE..\n");

/* Doing nothing at the moment.. */

break;
case PMS_PD_MGMT_STATUS_PMS_PERFORMANCE_DEGRADED:

printf("hasim :          received MGMT_STATUS PMS PERFORMANCE DEGRADED..\n");

/* Doing nothing at the moment.. */

break;
case PMS_PD_MGMT_STATUS_RND_ADDRESS_ADD:
case PMS_PD_MGMT_STATUS_RND_ADDRESS_DELETE:

if (pr->payload.data.mgmt_status.code == \
PMS_PD_MGMT_STATUS_RND_ADDRESS_ADD)
printf("hasim :          received MGMT_STATUS RND ADDRESS ADD..\n");
else
printf("hasim :          received MGMT_STATUS RND ADDRESS DELETE..\n");

info_get_fail = 0;

/* Get MGMT rnd address information.. */

ps.session.type = PMS_SR_CALL_RETURN_TIMED;
ps.session.info.crt.time = 0;
ps.payload.type = PMS_PD_MGMT_RND_ADDRESS_INFO_GET_EXECUTE;

for(i=0;i<16;i++)
{
ps.payload.data.mgmt_rnd_address_info_get_execute.index = i;

if (pms_send(&ps, &prs) == 0)
{
if (prs.payload.data.mgmt_rnd_address_info_get_status.err == \
PMS_PD_MGMT_RND_ADDRESS_INFO_GET_STATUS_ERR_NONE)
{
rnd_address_identifier[i] = \
prs.payload.data.mgmt_rnd_address_info_get_status.identifier;
strncpy(&rnd_address_node_ip_address[i][0], \
&prs.payload.data.mgmt_rnd_address_info_get_status.node_ip_address[0], 20);
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (Continued)

```
        strncpy(&rnd_address_drawer_ip_address[i][0], \
&prs.payload.data.mgmt_rnd_address_info_get_status.drawer_ip_address[0], 20);
        rnd_address_node_slot_number[i] = \
        prs.payload.data.mgmt_rnd_address_info_get_status.node_slot_number;
    }
    else
    {
        info_get_fail = 1;
    }
}
else
{
    info_get_fail = 1;
}
}

if (info_get_fail == 0)
{

    for(i=0;i<16;i++)
    {
        mdi.mgmt_info.rnd_address_cache[i].identifier = \
        rnd_address_identifier[i];
        strncpy(&mdi.mgmt_info.rnd_address_cache[i].node_ip_address[0], \
        &rnd_address_node_ip_address[i][0], 20);
        strncpy(&mdi.mgmt_info.rnd_address_cache[i].drawer_ip_address[0], \
        &rnd_address_drawer_ip_address[i][0], 20);
        mdi.mgmt_info.rnd_address_cache[i].node_slot_number = \
        rnd_address_node_slot_number[i];
    }

    mdi.mgmt_info.last_update = 0;
}

app_hasim_process();

break;
case PMS_PD_MGMT_STATUS_PMS_ADMIN_STATE_AV_RGOVA_DELAY:

    printf("hasim :          received MGMT_STATUS PMS ADMIN STATE AV RGOVA \
    DELAY..\n");
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
        /* Doing nothing at the moment.. */

        break;
    }

}

void
app_hasim_receive_node_rg0_status(struct pms_receive* pr)
{

    switch(pr->payload.data.node_rg0_status.code)
    {
        case PMS_PD_NODE_RG0_STATUS_STATE_ACTIVE:

            printf("hasim :          received NODE_RG0_STATUS STATE ACTIVE..\n");

            mdi.node_info.rg0_state_cache = PMS_PD_NODE_RG0_INFO_GET_STATUS_ACTIVE;
            mdi.node_info.last_update = 0;

            app_hasim_process();

            break;
        case PMS_PD_NODE_RG0_STATUS_STATE_OFFLINE:

            printf("hasim :          received NODE_RG0_STATUS STATE OFFLINE..\n");

            mdi.node_info.rg0_state_cache = PMS_PD_NODE_RG0_INFO_GET_STATUS_OFFLINE;
            mdi.node_info.last_update = 0;

            app_hasim_process();

            break;
        case PMS_PD_NODE_RG0_STATUS_ADMIN_STATE_FORCE_OFFLINE:

            printf("hasim :          received NODE_RG0_STATUS ADMIN STATE FORCE \
                OFFLINE ..\n");

            /* Doing nothing at the moment.. */

            break;
        case PMS_PD_NODE_RG0_STATUS_ADMIN_STATE_VOTE_ACTIVE:

            printf("hasim :          received NODE_RG0_STATUS ADMIN STATE VOTE \
                ACTIVE..\n");
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
    /* Doing nothing at the moment.. */

break;
case PMS_PD_NODE_RG0_STATUS_ADMIN_STATE_FORCE_ACTIVE:

    printf("hasim :      received NODE_RG0_STATUS ADMIN STATE FORCE \
ACTIVE..\n");

    /* Doing nothing at the moment.. */

break;
case PMS_PD_NODE_RG0_STATUS_APP_STATE_SET_FAULT:

    printf("hasim :      received NODE_RG0_STATUS APP STATE SET FAULT..\n");

    /* Doing nothing at the moment.. */

break;
case PMS_PD_NODE_RG0_STATUS_ADOPER_STATUSMASK_SET:

    printf("hasim :      received NODE_RG0_STATUS ADOPER STATUSMASK SET..\n");

    /* Doing nothing at the moment.. */

break;
}

}

void
app_hasim_receive_node_rg0_app_state_set_execute(struct pms_receive* pr)
{

    struct pms_send          ps;

    switch(pr->payload.data.node_rg0_app_state_set_execute.state)
    {
        case PMS_PD_NODE_RG0_APP_STATE_SET_EXECUTE_ACTIVE:

            printf("hasim :      received NODE_RG0_APP_STATE_SET_EXECUTE ACTIVE..\n");

            /* Do whatever, within pr->session.info.crt.time if possible.. */
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
    /* Send return message indicating successful reception.. */

    ps.session.type = PMS_SR_RETURN;
    ps.session.info.r.return_identifier = \
        pr->session.info.crt.call_identifier;
    ps.session.info.r.return_priority = pr->session.info.crt.return_priority;

    ps.payload.type = PMS_PD_NODE_RG0_APP_STATE_SET_STATUS;
    ps.payload.data.node_rg0_app_state_set_status.err = \
        PMS_PD_NODE_RG0_APP_STATE_SET_STATUS_SUCCESS;

    if (pms_send(&ps, 0) != 0)
    {
    }

    break;
    case PMS_PD_NODE_RG0_APP_STATE_SET_EXECUTE_OFFLINE:

        printf("hasim :      received NODE_RG0_APP_STATE_SET_EXECUTE OFFLINE\
            ..\n");

        /* Do whatever, within pr->session.info.crt.time if possible.. */

        ps.session.type = PMS_SR_RETURN;
        ps.session.info.r.return_identifier = \
            pr->session.info.crt.call_identifier;
        ps.session.info.r.return_priority = pr->session.info.crt.return_priority;

        ps.payload.type = PMS_PD_NODE_RG0_APP_STATE_SET_STATUS;
        ps.payload.data.node_rg0_app_state_set_status.err = \
            PMS_PD_NODE_RG0_APP_STATE_SET_STATUS_SUCCESS;

        if (pms_send(&ps, 0) != 0)
        {
        }

        break;
    };
}
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
void
app_hasim_receive_rnd_status(struct pms_receive* pr)
{

    printf("hasim :          rs.identifier=%.8X\n", \
           pr->payload.data.rnd_status.identifier);

    switch(pr->payload.data.rnd_status.code)
    {
    case PMS_PD_RND_STATUS_VIEW_NODE_REACHABLE_DRAWER_REACHABLE:

        printf("hasim :          received RND_STATUS NODE REACHABLE DRAWER \
               REACHABLE..\n");

        /* Doing nothing at the moment.. */

        break;
    case PMS_PD_RND_STATUS_VIEW_NODE_REACHABLE_DRAWER_UNREACHABLE:

        printf("hasim :          received RND_STATUS NODE REACHABLE DRAWER\
               UNREACHABLE..\n");

        /* Doing nothing at the moment.. */

        break;
    case PMS_PD_RND_STATUS_VIEW_NODE_UNREACHABLE_DRAWER_REACHABLE:

        printf("hasim :          received RND_STATUS NODE UNREACHABLE DRAWER\
               REACHABLE..\n");

        /* Doing nothing at the moment.. */

        break;
    case PMS_PD_RND_STATUS_VIEW_NODE_UNREACHABLE_DRAWER_UNREACHABLE:

        printf("hasim :          received RND_STATUS NODE UNREACHABLE DRAWER\
               UNREACHABLE..\n");

        /* Doing nothing at the moment.. */

        break;
    case PMS_PD_RND_STATUS_ADOPER_FORCE_UNAVAILABLE:

        printf("hasim :          received RND_STATUS ADOPER FORCE UNAVAILABLE..\n");
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
    /* Doing nothing at the moment.. */

    break;
    case PMS_PD_RND_STATUS_ADOPER_VOTE_AVAILABLE:

        printf("hasim :          received RND_STATUS ADOPER VOTE AVAILABLE..\n");

        /* Doing nothing at the moment.. */

    break;
    case PMS_PD_RND_STATUS_ADOPER_FORCE_AVAILABLE:

        printf("hasim :          received RND_STATUS ADOPER FORCE AVAILABLE..\n");

        /* Doing nothing at the moment.. */

    break;
    case PMS_PD_RND_STATUS_ADOPER_STATUSMASK_SET:

        printf("hasim :          received RND_STATUS ADOPER STATUSMASK SET..\n");

        /* Doing nothing at the moment.. */

    break;
    case PMS_PD_RND_STATUS_STATE_UNAVAILABLE:

        printf("hasim :          received RND_STATUS STATE UNAVAILABLE..\n");

        /* Doing nothing at the moment.. */

    break;
    case PMS_PD_RND_STATUS_STATE_AVAILABLE:

        printf("hasim :          received RND_STATUS STATE AVAILABLE..\n");

        /* Doing nothing at the moment.. */

    break;
}

void
app_hasim_receive_rnd_md0_status(struct pms_receive* pr)
{
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (Continued)

```
printf("hasim :          rms.identifier=%.8X\n", \
      pr->payload.data.rnd_md0_status.identifier);

switch(pr->payload.data.rnd_md0_status.code)
{
case PMS_PD_RND_MD0_STATUS_ADOPER_CONFIG_MAINTENANCE:

    printf("hasim :          received RND_MD0_STATUS ADOPER CONFIG \
          MAINTENANCE..\n");

    /* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_CONFIG_OPERATIONAL:

    printf("hasim :          received RND_MD0_STATUS ADOPER CONFIG \
          OPERATIONAL..\n");

    /* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_GRACEFUL_REBOOT:

    printf("hasim :          received RND_MD0_STATUS ADOPER GRACEFUL REBOOT..\n");

    /* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_STATUSMASK_SET:

    printf("hasim :          received RND_MD0_STATUS ADOPER STATUSMASK SET..\n");

    /* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_RECOVERY_PC:

    printf("hasim :          received RND_MD0_STATUS ADOPER RECOVERY PC..\n");

    /* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_RECOVERY_RST:
```


CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
printf("hasim :          received RND_MD0_STATUS ADOPER RECOVERY RST..\n");

/* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_RECOVERY_RSTPC:

printf("hasim :          received RND_MD0_STATUS ADOPER RECOVERY RSTPC..\n");

/* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_RECOVERY_PD:
printf("hasim :          received RND_MD0_STATUS ADOPER RECOVERY PD..\n");

/* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_RECOVERY_RB:

printf("hasim :          received RND_MD0_STATUS ADOPER RECOVERY RB..\n");

/* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_RECOVERYAUTOMODE_SET:

printf("hasim :          received RND_MD0_STATUS ADOPER RECOVERYAUTOMODE\
      SET..\n");

/* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_ADOPER_SCDM_TIMEOUT:

printf("hasim :          received RND_MD0_STATUS ADOPER SCDM TIMEOUT..\n");

/* Doing nothing at the moment.. */

break;
case PMS_PD_RND_MD0_STATUS_CONFIG_MAINTENANCE:

printf("hasim :          received RND_MD0_STATUS CONFIG MAINTENANCE..\n");

/* Doing nothing at the moment.. */

break;
```

CODE EXAMPLE 7-3 PMS Client Asynchronous Message Handling (*Continued*)

```
case PMS_PD_RND_MD0_STATUS_CONFIG_OPERATIONAL:

    printf("hasim :      received RND_MD0_STATUS CONFIG OPERATIONAL...\n");

    /* Doing nothing at the moment.. */

    break;
}

}
```

The following example shows a PMS client's scheduling.

CODE EXAMPLE 7-4 PMS Client Scheduling Example

```
void
app_hasim_receive_time_status(void)
{

    int                i;
    mdi.mgmt_info.last_update += HASIM_CHECK_INTERVAL;
    mdi.mgmt_info.last_sync_check += HASIM_CHECK_INTERVAL;

    mdi.node_info.last_update += HASIM_CHECK_INTERVAL;
    mdi.node_info.last_sync_check += HASIM_CHECK_INTERVAL;

    for(i=0;i<16;i++)
    {
        mdi.rnd_info[i].last_update += HASIM_CHECK_INTERVAL;
        mdi.rnd_info[i].last_sync_check += HASIM_CHECK_INTERVAL;
    }

    app_hasim_process();

}

void
app_hasim_process(void)
{
```

CODE EXAMPLE 7-4 PMS Client Scheduling Example (*Continued*)

```
/* Run state machine sub-groupings.. */  
  
app_hasim_user_process();  
  
app_hasim_mgmt_process();  
  
app_hasim_node_process();  
  
app_hasim_rnd_process();  
  
}
```

The following example shows the PMS client's user management interface.

CODE EXAMPLE 7-5 PMS Client User and Management Interface

```
void  
app_hasim_user_process(void)  
{  
  
    struct pms_receive    pr;  
  
    int                   i;  
  
  
    /* PMS View check */  
  
    /* Periodically attempt to connect if unreachable. Return to initial  
       state variable settings on reachable to unreachable transition.. */  
  
    if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_UNREACHABLE)  
    {  
        if (pms_connect(PMS_SERVER_PORT_NUMBER_DEFAULT) != 0)  
        {  
        }  
        else  
        {  
            mdi.user_info.pms_view = HASIM_USER_PMS_VIEW_REACHABLE;  
            mdi.user_info.view_cache = PMS_PD_USER_STATUS_PMS_REACHABLE;  
        }  
    }  
}
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
else /* mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE */
{
    if (mdi.user_info.view_cache == PMS_PD_USER_STATUS_PMS_UNREACHABLE)
    {

        /* RND */

        for(i=0;i<16;i++)
        {

            mdi.rnd_info[i].cache_state = HASIM_RND_CACHE_INVALID;
            mdi.rnd_info[i].last_update = HASIM_CHECK_INVALID_INTERVAL;

            if ((mdi.rnd_info[i].receive_state & HASIM_RND_RECEIVE_REGISTERED) != 0)
            {
                pr.session.type = PMS_SR_CALL_NO_RETURN;
                pr.payload.type = PMS_PD_RND_STATUS;
                pr.payload.data.rnd_status.identifier = \
                    mdi.mgmt_info.rnd_address_identifier[i];
                pms_receive(&pr, 0, 0);
                mdi.rnd_info[i].receive_state &= !HASIM_RND_RECEIVE_REGISTERED;
            }

            if ((mdi.rnd_info[i].receive_state & HASIM_RND_MD0_RECEIVE_REGISTERED)\
                != 0)
            {
                pr.session.type = PMS_SR_CALL_NO_RETURN;
                pr.payload.type = PMS_PD_RND_MD0_STATUS;
                pr.payload.data.rnd_status.identifier = \
                    mdi.mgmt_info.rnd_address_identifier[i];
                pms_receive(&pr, 0, 0);
                mdi.rnd_info[i].receive_state &= !HASIM_RND_MD0_RECEIVE_REGISTERED;
            }

        }

        /* NODE */

        mdi.node_info.cache_state = HASIM_NODE_CACHE_INVALID;
        mdi.node_info.last_update = HASIM_CHECK_INVALID_INTERVAL;

        if (mdi.node_info.service_state == HASIM_NODE_SERVICE_STATE_ACTIVE)
        {
            mdi.node_info.service_state = HASIM_NODE_SERVICE_STATE_OFFLINE;
        }
    }
}
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (*Continued*)

```
if (mdi.node_info.rg0_app_name_state ==\
    HASIM_NODE_RG0_APP_NAME_REGISTERED)
{
mdi.node_info.rg0_app_name_state = HASIM_NODE_RG0_APP_NAME_UNREGISTERED;
}

if ((mdi.node_info.receive_state &\
    HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED) != 0)
{
pr.session.type = PMS_SR_CALL_NO_RETURN;
pr.payload.type = PMS_PD_NODE_RG0_STATUS;
pms_receive(&pr, 0, 0);
mdi.node_info.receive_state &=\
    !HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED;
}

if ((mdi.node_info.receive_state &\
    HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED) != 0)
{
pr.session.type = PMS_SR_CALL_RETURN_TIMED;
pr.payload.type = PMS_PD_NODE_RG0_APP_STATE_SET_EXECUTE;
pms_receive(&pr, 0, 0);
mdi.node_info.receive_state &= \
    !HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED;
}

/* MGMT */

mdi.mgmt_info.cache_state = HASIM_MGMT_CACHE_INVALID;
mdi.mgmt_info.last_update = HASIM_CHECK_INVALID_INTERVAL;

for(i=0;i<16;i++)
    mdi.mgmt_info.rnd_address_identifier[i] = -1;

if (mdi.mgmt_info.rnd_address_state == HASIM_MGMT_RND_ADDRESS_VERIFIED)
{
mdi.mgmt_info.rnd_address_state = HASIM_MGMT_RND_ADDRESS_UNVERIFIED;
}

if (mdi.mgmt_info.pms_state == HASIM_MGMT_PMS_STATE_AVAILABLE)
{
mdi.mgmt_info.pms_state = HASIM_MGMT_PMS_STATE_UNAVAILABLE;
}

if (mdi.mgmt_info.receive_state == HASIM_MGMT_RECEIVE_REGISTERED)
{
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
pr.session.type = PMS_SR_CALL_NO_RETURN;
pr.payload.type = PMS_PD_MGMT_STATUS;
pms_receive(&pr, 0, 1);

mdi.mgmt_info.receive_state = HASIM_MGMT_RECEIVE_UNREGISTERED;
}

/* USER */

if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
{
pr.session.type = PMS_SR_CALL_NO_RETURN;
pr.payload.type = PMS_PD_USER_STATUS;
pms_receive(&pr, 0, 0);

mdi.user_info.receive_state = HASIM_USER_RECEIVE_UNREGISTERED;
}

if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
{
pms_disconnect();

mdi.user_info.pms_view = HASIM_USER_PMS_VIEW_UNREACHABLE;
}
}

}

/* Receive Check */

/* If USER messages are not receive registered, attempt to register if PMS
is reachable.. */

if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
{
if (mdi.user_info.receive_state != HASIM_USER_RECEIVE_REGISTERED)
{
pr.session.type = PMS_SR_CALL_NO_RETURN;
pr.payload.type = PMS_PD_USER_STATUS;
if (pms_receive(&pr, app_hasim_receive_post, 0) != -1)
mdi.user_info.receive_state = HASIM_USER_RECEIVE_REGISTERED;
}
}
}
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (*Continued*)

```
}

void
app_hasim_mgmt_process(void)
{
    struct pms_send      ps;
    struct pms_receive   pr;

    int                  info_get_fail;

    int                  match[HASIM_RND_ADDRESS_AUDIT_ENTRIES];

    int                  mgmt_state;
    int                  rnd_address_identifier[16];
    char                 rnd_address_node_ip_address[16][20];
    char                 rnd_address_drawer_ip_address[16][20];
    int                  rnd_address_node_slot_number[16];

    int                  i, j;

    /* Receive Check */

    /* If MGMT messages are not receive registered, attempt to register if PMS
       is reachable and USER receive messages are registered.  If registration
       is successful, force an initial cache update.. */

    if (mdi.mgmt_info.receive_state != HASIM_MGMT_RECEIVE_REGISTERED)
    {
        if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
        {
            if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
            {
                pr.session.type = PMS_SR_CALL_NO_RETURN;
                pr.payload.type = PMS_PD_MGMT_STATUS;
                if (pms_receive(&pr, app_hasim_receive_post, 1) != -1)
                    mdi.mgmt_info.receive_state = HASIM_MGMT_RECEIVE_REGISTERED;

                /* Force an info_get immediately after registering.. */
                mdi.mgmt_info.last_sync_check = HASIM_SYNCCHECK_INTERVAL;
            }
        }
    }
}
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
/* PMS State check */

/* Process PMS state transitions. On an available to unavailable transition
return to pre-NODE and RND operational state variable settings.. */

if (mdi.mgmt_info.pms_state == HASIM_MGMT_PMS_STATE_UNAVAILABLE)
{
    if (mdi.mgmt_info.cache_state != HASIM_MGMT_CACHE_INVALID)
    {
        if (mdi.mgmt_info.mgmt_state_cache != \
            PMS_PD_MGMT_INFO_GET_STATUS_UNAVAILABLE)
        {
            mdi.mgmt_info.pms_state = HASIM_MGMT_PMS_STATE_AVAILABLE;
        }
    }
}
else /* mdi.mgmt_info.pms_state == HASIM_MGMT_PMS_STATE_AVAILABLE */
{
    if (mdi.mgmt_info.cache_state != HASIM_MGMT_CACHE_INVALID)
    {
        if (mdi.mgmt_info.mgmt_state_cache == \
            PMS_PD_MGMT_INFO_GET_STATUS_UNAVAILABLE)
        {

            /* RND */

            for(i=0;i<16;i++)
            {

                if ((mdi.rnd_info[i].receive_state & \
                    HASIM_RND_RECEIVE_REGISTERED) != 0)
                {
                    pr.session.type = PMS_SR_CALL_NO_RETURN;
                    pr.payload.type = PMS_PD_RND_STATUS;
                    pr.payload.data.rnd_status.identifier = \
                        mdi.mgmt_info.rnd_address_identifier[i];
                    pms_receive(&pr, 0, 0);
                    mdi.rnd_info[i].receive_state &= !HASIM_RND_RECEIVE_REGISTERED;
                }

                if ((mdi.rnd_info[i].receive_state & HASIM_RND_MD0_RECEIVE_REGISTERED) \
                    != 0)
                {
                    pr.session.type = PMS_SR_CALL_NO_RETURN;
                }
            }
        }
    }
}
```


CODE EXAMPLE 7-5 PMS Client User and Management Interface (*Continued*)

```
        pr.payload.type = PMS_PD_RND_MD0_STATUS;
        pr.payload.data.rnd_status.identifier = \
            mdi.mgmt_info.rnd_address_identifier[i];
        pms_receive(&pr, 0, 0);
        mdi.rnd_info[i].receive_state &= !HASIM_RND_MD0_RECEIVE_REGISTERED;
    }

}

/* NODE */

if (mdi.node_info.service_state == HASIM_NODE_SERVICE_STATE_ACTIVE)
{
    mdi.node_info.service_state = HASIM_NODE_SERVICE_STATE_OFFLINE;
}

if (mdi.node_info.rg0_app_name_state == \
    HASIM_NODE_RG0_APP_NAME_REGISTERED)
{
    mdi.node_info.rg0_app_name_state = HASIM_NODE_RG0_APP_NAME_UNREGISTERED;
}

if ((mdi.node_info.receive_state & \
    HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED) != 0)
{
    pr.session.type = PMS_SR_CALL_NO_RETURN;
    pr.payload.type = PMS_PD_NODE_RG0_STATUS;
    pms_receive(&pr, 0, 0);
    mdi.node_info.receive_state &= \
        !HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED;
}

if ((mdi.node_info.receive_state & \
    HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED) != 0)
{
    pr.session.type = PMS_SR_CALL_RETURN_TIMED;
    pr.payload.type = PMS_PD_NODE_RG0_APP_STATE_SET_EXECUTE;
    pms_receive(&pr, 0, 0);
    mdi.node_info.receive_state &= \
        !HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED;
}

mdi.mgmt_info.pms_state = HASIM_MGMT_PMS_STATE_UNAVAILABLE;
}
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
    }
}

/* RND Address Check */

/* Check once at startup if the RND address pairs currently in the list
are the same as this control application's defaults. If not, remove
any that differ and add any that are missing. This is a bit contrived
to demonstrate interaction via the address list messages. No point
in starting processing if cache is invalid and PMS is not reachable
and USER registration is not completed.. */

if (mdi.mgmt_info.rnd_address_state != HASIM_MGMT_RND_ADDRESS_VERIFIED)
{
    if (mdi.mgmt_info.cache_state != HASIM_MGMT_CACHE_INVALID)
    {
        if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
        {
            if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
            {

                match[0] = 0;
                match[1] = 0;

                /* Search RND address list for entries not in the app's verify list.. */

                for(i=0;i<16;i++)
                {
                    if (mdi.mgmt_info.rnd_address_cache[i].identifier != -1)
                    {
                        for(j=0;j<HASIM_RND_ADDRESS_AUDIT_ENTRYS;j++)
                        {
                            if (match[j] == 0)
                            {
                                /* Use strcmp() for the moment. Use sockaddr_in when \
I get around to it.. */
                                if \
                                (strcmp(&mdi.mgmt_info.rnd_address_cache[i].node_ip_address[0], \
&mdi.rnd_address[j].node_ip_address[0]) == 0)
                                {
                                    if\
                                    (strcmp(&mdi.mgmt_info.rnd_address_cache[i].drawer_ip_address[0], \
&mdi.rnd_address[j].drawer_ip_address[0]) == 0)
                                    {
                                        if (mdi.mgmt_info.rnd_address_cache[i].node_slot_number == \
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
        mdi.rnd_address[j].node_slot_number)
        {
            match[j] = 1;

            break;
        }
    }
}

/* Delete entries not in the app's verify list.. */

if (j == HASIM_RND_ADDRESS_AUDIT_ENTRIES)
{
    ps.session.type = PMS_SR_CALL_RETURN_TIMED;
    ps.session.info.crt.time = 0;
    ps.payload.type = PMS_PD_MGMT_RND_ADDRESS_DELETE_EXECUTE;

    ps.payload.data.mgmt_rnd_address_delete_execute.identifier = \
        mdi.mgmt_info.rnd_address_cache[i].identifier;

    if (pms_send(&ps, &pr) == 0)
    {
        {
            if (pr.payload.data.mgmt_rnd_address_delete_status.err == \
                PMS_PD_MGMT_RND_ADDRESS_DELETE_STATUS_ERR_NONE)
            {
            }
        }
    }
}

}

/* Add any missing entries.. */

for(i=0;i<HASIM_RND_ADDRESS_AUDIT_ENTRIES;i++)
{
    if (match[i] == 0)
    {
        ps.session.type = PMS_SR_CALL_RETURN_TIMED;
        ps.session.info.crt.time = 0;
        ps.payload.type = PMS_PD_MGMT_RND_ADDRESS_ADD_EXECUTE;

        strncpy(&ps.payload.data.mgmt_rnd_address_add_execute.node_ip_address[0], \
            &mdi.rnd_address[i].node_ip_address[0], 20);
        strncpy(&ps.payload.data.mgmt_rnd_address_add_execute.drawer_ip_address[0], \
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
&mdi.rnd_address[i].drawer_ip_address[0], 20);
ps.payload.data.mgmt_rnd_address_add_execute.node_slot_number = \
mdi.rnd_address[i].node_slot_number;

    if (pms_send(&ps, &pr) == 0)
    {
        if (pr.payload.data.mgmt_rnd_address_add_status.err == \
            PMS_PD_MGMT_RND_ADDRESS_ADD_STATUS_ERR_NONE)
        {
            }
        }
    }
}

    mdi.mgmt_info.rnd_address_state = HASIM_MGMT_RND_ADDRESS_VERIFIED;
}
}
}

/* RND Address Identifier check */

/* Process RND address identifier transitions. On in-use to not-in-use
transitions, return state variables to pre-RND initialized state for that
identifier. Check whether any list entries have been deleted and re-added
since last processing and do an available->unavailable->available
transition.. */

for(i=0;i<16;i++)
{

    if (mdi.mgmt_info.rnd_address_identifier[i] == -1)
    {

        if (mdi.mgmt_info.cache_state != HASIM_MGMT_CACHE_INVALID)
        {
            if (mdi.mgmt_info.rnd_address_cache[i].identifier != -1)
            {
                mdi.mgmt_info.rnd_address_identifier[i] =\
                    mdi.mgmt_info.rnd_address_cache[i].identifier;
            }
        }

    }

    else /* mdi.mgmt_info.rnd_address_identifier[i] != -1 */
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (*Continued*)

```
{

if (mdi.mgmt_info.cache_state != HASIM_MGMT_CACHE_INVALID)
{

if (mdi.mgmt_info.rnd_address_cache[i].identifier == -1)
{

/* RND */

if ((mdi.rnd_info[i].receive_state & HASIM_RND_RECEIVE_REGISTERED)\
    != 0)
{
pr.session.type = PMS_SR_CALL_NO_RETURN;
pr.payload.type = PMS_PD_RND_STATUS;
pr.payload.data.rnd_status.identifier = \
    mdi.mgmt_info.rnd_address_identifier[i];
pms_receive(&pr, 0, 0);
mdi.rnd_info[i].receive_state &= !HASIM_RND_RECEIVE_REGISTERED;
}

if ((mdi.rnd_info[i].receive_state & HASIM_RND_MD0_RECEIVE_REGISTERED)\
    != 0)
{
pr.session.type = PMS_SR_CALL_NO_RETURN;
pr.payload.type = PMS_PD_RND_MD0_STATUS;
pr.payload.data.rnd_status.identifier = \
    mdi.mgmt_info.rnd_address_identifier[i];
pms_receive(&pr, 0, 0);
mdi.rnd_info[i].receive_state &= !HASIM_RND_MD0_RECEIVE_REGISTERED;
}

mdi.mgmt_info.rnd_address_identifier[i] = -1;
}
else
{
if (mdi.mgmt_info.rnd_address_identifier[i] != \
    mdi.mgmt_info.rnd_address_cache[i].identifier)
{
/* RND */

if ((mdi.rnd_info[i].receive_state & HASIM_RND_RECEIVE_REGISTERED)\
    != 0)
{
pr.session.type = PMS_SR_CALL_NO_RETURN;
pr.payload.type = PMS_PD_RND_STATUS;
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
pr.payload.data.rnd_status.identifier = \  
    mdi.mgmt_info.rnd_address_identifier[i];  
pms_receive(&pr, 0, 0);  
mdi.rnd_info[i].receive_state &= !HASIM_RND_RECEIVE_REGISTERED;  
}  
  
if ((mdi.rnd_info[i].receive_state & \  
    HASIM_RND_MD0_RECEIVE_REGISTERED) != 0)  
{  
    pr.session.type = PMS_SR_CALL_NO_RETURN;  
    pr.payload.type = PMS_PD_RND_MD0_STATUS;  
    pr.payload.data.rnd_status.identifier = \  
        mdi.mgmt_info.rnd_address_identifier[i];  
    pms_receive(&pr, 0, 0);  
    mdi.rnd_info[i].receive_state &= !HASIM_RND_MD0_RECEIVE_REGISTERED;  
}  
  
    mdi.mgmt_info.rnd_address_identifier[i] = \  
        mdi.mgmt_info.rnd_address_cache[i].identifier;  
}  
}  
  
}  
  
}  
  
}  
  
/* Sync Check */  
  
/* Policy: Sync update checked every SYNCHECK_INTERVAL seconds.. */  
if (mdi.mgmt_info.last_sync_check > HASIM_SYNCHECK_INTERVAL)  
{  
    /* Policy: Don't attempt a sync update if any async partial updates have  
        been received within SYNCHECK_INTERVAL.. */  
    if (mdi.mgmt_info.last_update > HASIM_SYNCHECK_INTERVAL)  
    {  
        /* Policy: Don't attempt a sync update if registration for async  
            updates have not succeeded.. */  
        if (mdi.mgmt_info.receive_state == HASIM_MGMT_RECEIVE_REGISTERED)  
        {  
            if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)  
            {  
                if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)  
                {  
                    mdi.mgmt_info.last_sync_check = 0;  
                }  
            }  
        }  
    }  
}
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (*Continued*)

```
info_get_fail = 0;

/* Get MGMT base information.. */
ps.session.type = PMS_SR_CALL_RETURN_TIMED;
ps.session.info.crt.time = 0;
ps.payload.type = PMS_PD_MGMT_INFO_GET_EXECUTE;

if (pms_send(&ps, &pr) == 0)
{
    if (pr.payload.data.mgmt_info_get_status.err == \
        PMS_PD_MGMT_INFO_GET_STATUS_SUCCESS)
    {
        mgmt_state = pr.payload.data.mgmt_info_get_status.state;
    }
    else
    {
        info_get_fail = 1;
    }
}
else
{
    info_get_fail = 1;
}

/* Get MGMT rnd address information.. */

ps.session.type = PMS_SR_CALL_RETURN_TIMED;
ps.session.info.crt.time = 0;
ps.payload.type = PMS_PD_MGMT_RND_ADDRESS_INFO_GET_EXECUTE;

for(i=0;i<16;i++)
{
    ps.payload.data.mgmt_rnd_address_info_get_execute.index = i;

    if (pms_send(&ps, &pr) == 0)
    {
        if (pr.payload.data.mgmt_rnd_address_info_get_status.err == \
            PMS_PD_MGMT_RND_ADDRESS_INFO_GET_STATUS_ERR_NONE)
        {
            rnd_address_identifer[i] = \
                pr.payload.data.mgmt_rnd_address_info_get_status.identifer;
            strncpy(&rnd_address_node_ip_address[i][0], \
                &pr.payload.data.mgmt_rnd_address_info_get_status.node_ip_address[0], 20);
            strncpy(&rnd_address_drawer_ip_address[i][0], \
                &pr.payload.data.mgmt_rnd_address_info_get_status.drawer_ip_address[0], 20);
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
        rnd_address_node_slot_number[i] = \
pr.payload.data.mgmt_rnd_address_info_get_status.node_slot_number;
    }
    else
    {
        info_get_fail = 1;
    }
}
else
{
    info_get_fail = 1;
}
}

/* Only mark MGMT update as successful if all pieces of data
   were received successfully.. */

if (info_get_fail == 0)
{
    mdi.mgmt_info.mgmt_state_cache = mgmt_state;

    for(i=0;i<16;i++)
    {
        mdi.mgmt_info.rnd_address_cache[i].identifier = \
            rnd_address_identifier[i];
        strncpy(&mdi.mgmt_info.rnd_address_cache[i].node_ip_address[0], \
            &rnd_address_node_ip_address[i][0], 20);
        strncpy(&mdi.mgmt_info.rnd_address_cache[i].drawer_ip_address[0], \
            &rnd_address_drawer_ip_address[i][0], 20);
        mdi.mgmt_info.rnd_address_cache[i].node_slot_number = \
            rnd_address_node_slot_number[i];
    }

    mdi.mgmt_info.last_update = 0;
}
}
}
}
else
{
    mdi.mgmt_info.last_sync_check = 0;
}
}
```


CODE EXAMPLE 7-5 PMS Client User and Management Interface (*Continued*)

```
/* Validity Check */

/* Process cache state validity transitions. The policy is on a MGMT cache
   transition to invalid, return state variables to initial configuration.. */

if(mdi.mgmt_info.last_update < HASIM_CHECK_VALID_INTERVAL)
{
    if (mdi.mgmt_info.cache_state != HASIM_MGMT_CACHE_VALID)
        mdi.mgmt_info.cache_state = HASIM_MGMT_CACHE_VALID;
}
else if((mdi.mgmt_info.last_update >= HASIM_CHECK_VALID_INTERVAL && \
        mdi.mgmt_info.last_update < HASIM_CHECK_INVALID_INTERVAL))
{
    if (mdi.mgmt_info.cache_state == HASIM_MGMT_CACHE_VALID)
        mdi.mgmt_info.cache_state = HASIM_MGMT_CACHE_OLD;
}
else if(mdi.mgmt_info.last_update >= HASIM_CHECK_INVALID_INTERVAL)
{
    if (mdi.mgmt_info.cache_state == HASIM_MGMT_CACHE_OLD)
    {

        /* RND */

        for(i=0;i<16;i++)
        {

            mdi.rnd_info[i].cache_state = HASIM_RND_CACHE_INVALID;
            mdi.rnd_info[i].last_update = HASIM_CHECK_INVALID_INTERVAL;

            if ((mdi.rnd_info[i].receive_state & HASIM_RND_RECEIVE_REGISTERED) != 0)
            {
                pr.session.type = PMS_SR_CALL_NO_RETURN;
                pr.payload.type = PMS_PD_RND_STATUS;
                pr.payload.data.rnd_status.identifier = \
                    mdi.mgmt_info.rnd_address_identifier[i];
                pms_receive(&pr, 0, 0);
                mdi.rnd_info[i].receive_state &= !HASIM_RND_RECEIVE_REGISTERED;
            }

            if ((mdi.rnd_info[i].receive_state & HASIM_RND_MD0_RECEIVE_REGISTERED)\
                != 0)
            {
                pr.session.type = PMS_SR_CALL_NO_RETURN;
                pr.payload.type = PMS_PD_RND_MD0_STATUS;
                pr.payload.data.rnd_status.identifier = \
                    mdi.mgmt_info.rnd_address_identifier[i];
                pms_receive(&pr, 0, 0);
            }
        }
    }
}
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
        mdi.rnd_info[i].receive_state &= !HASIM_RND_MD0_RECEIVE_REGISTERED;
    }

}

/* NODE*/

mdi.node_info.cache_state = HASIM_NODE_CACHE_INVALID;
mdi.node_info.last_update = HASIM_CHECK_INVALID_INTERVAL;

if (mdi.node_info.service_state == HASIM_NODE_SERVICE_STATE_ACTIVE)
{
    mdi.node_info.service_state = HASIM_NODE_SERVICE_STATE_OFFLINE;
}

if (mdi.node_info.rg0_app_name_state ==\
    HASIM_NODE_RG0_APP_NAME_REGISTERED)
{
    mdi.node_info.rg0_app_name_state = HASIM_NODE_RG0_APP_NAME_UNREGISTERED;
}

if ((mdi.node_info.receive_state &\
    HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED) != 0)
{
    pr.session.type = PMS_SR_CALL_NO_RETURN;
    pr.payload.type = PMS_PD_NODE_RG0_STATUS;
    pms_receive(&pr, 0, 0);
    mdi.node_info.receive_state &=\
        !HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED;
}

if ((mdi.node_info.receive_state &\
    HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED) != 0)
{
    pr.session.type = PMS_SR_CALL_RETURN_TIMED;
    pr.payload.type = PMS_PD_NODE_RG0_APP_STATE_SET_EXECUTE;
    pms_receive(&pr, 0, 0);
    mdi.node_info.receive_state &=\
        !HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED;
}

/* MGMT */

mdi.mgmt_info.cache_state = HASIM_MGMT_CACHE_INVALID;
```

CODE EXAMPLE 7-5 PMS Client User and Management Interface (Continued)

```
for(i=0;i<16;i++)
    mdi.mgmt_info.rnd_address_identifier[i] = -1;

if (mdi.mgmt_info.rnd_address_state == HASIM_MGMT_RND_ADDRESS_VERIFIED)
{
    mdi.mgmt_info.rnd_address_state = HASIM_MGMT_RND_ADDRESS_UNVERIFIED;
}

if (mdi.mgmt_info.pms_state == HASIM_MGMT_PMS_STATE_AVAILABLE)
{
    mdi.mgmt_info.pms_state = HASIM_MGMT_PMS_STATE_UNAVAILABLE;
}

if (mdi.mgmt_info.receive_state == HASIM_MGMT_RECEIVE_REGISTERED)
{
    pr.session.type = PMS_SR_CALL_NO_RETURN;
    pr.payload.type = PMS_PD_MGMT_STATUS;
    pms_receive(&pr, 0, 1);

    mdi.mgmt_info.receive_state = HASIM_MGMT_RECEIVE_UNREGISTERED;
}

/* USER */

if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
{
    pr.session.type = PMS_SR_CALL_NO_RETURN;
    pr.payload.type = PMS_PD_USER_STATUS;
    pms_receive(&pr, 0, 0);

    mdi.user_info.receive_state = HASIM_USER_RECEIVE_UNREGISTERED;
}

if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
{
    pms_disconnect();

    mdi.user_info.pms_view = HASIM_USER_PMS_VIEW_UNREACHABLE;
}
}
}
```

The following example shows the PMS client node interface.

CODE EXAMPLE 7-6 PMS Client Node Interface

```
void
app_hasim_node_process(void)
{

    struct pms_send      ps;
    struct pms_receive   pr;

    int                  info_get_fail;

    int                  rg0_state;

    int                  i;

    /* Receive Check */

    /* If NODE messages are not receive registered, attempt to register them if PMS
       is in the available state and reachable, and if USER receive messages are
       registered.  If registration is successful, force an initial cache
       update.. */

    if (mdi.node_info.receive_state != HASIM_NODE_GROUP_RECEIVE_REGISTERED)
    {
        if (mdi.mgmt_info.pms_state == HASIM_MGMT_PMS_STATE_AVAILABLE)
        {
            if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
            {
                if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
                {

                    if ((mdi.node_info.receive_state & \
                        HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED) == 0)
                    {
                        pr.session.type = PMS_SR_CALL_NO_RETURN;
                        pr.payload.type = PMS_PD_NODE_RG0_STATUS;
                        if (pms_receive(&pr, app_hasim_receive_post, 0) != -1)
                            mdi.node_info.receive_state |= \
                                HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED;
                    }

                    if ((mdi.node_info.receive_state & \
                        HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED) == 0)
                    {
                        pr.session.type = PMS_SR_CALL_RETURN_TIMED;
                        pr.session.info.crt.time = 50;
                    }
                }
            }
        }
    }
}
```

CODE EXAMPLE 7-6 PMS Client Node Interface (Continued)

```
        pr.payload.type = PMS_PD_NODE_RG0_APP_STATE_SET_EXECUTE;
        if (pms_receive(&pr, app_hasim_receive_post, 0) != -1)
            mdi.node_info.receive_state |= \
                HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED;
    }

    /* Force an info_get immediately after registering.. */
    mdi.node_info.last_sync_check = HASIM_SYNCCHECK_INTERVAL;
}
}
}

/* Name Check */

/* If this application's name is not registered, register it if PMS is
available and reachable, and if USER registration is complete.. */

if (mdi.node_info.rg0_app_name_state != HASIM_NODE_RG0_APP_NAME_REGISTERED)
{
    if (mdi.mgmt_info.pms_state == HASIM_MGMT_PMS_STATE_AVAILABLE)
    {
        if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
        {
            if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
            {

                /* Set NODE RG0 application name.. */

                ps.session.type = PMS_SR_CALL_RETURN_TIMED;
                ps.session.info.crt.time = 0;
                ps.payload.type = PMS_PD_NODE_RG0_APP_NAME_EXECUTE;
                strcpy(&ps.payload.data.node_rg0_app_name_execute.name[0], \
                    "hasim");
                ps.payload.data.node_rg0_app_name_execute.command = \
                    PMS_PD_NODE_RG0_APP_NAME_EXECUTE_ADD;

                if (pms_send(&ps, &pr) == 0)
                {
                    if (pr.payload.data.node_rg0_app_name_status.err == \
                        PMS_PD_NODE_RG0_APP_NAME_STATUS_ERR_NONE)
                    {
                        mdi.node_info.rg0_app_name_state = \
                            HASIM_NODE_RG0_APP_NAME_REGISTERED;
                    }
                }
            }
        }
    }
}
```

CODE EXAMPLE 7-6 PMS Client Node Interface (Continued)

```
    }
  }
}

/* Service State check */

/* Process application service state transitions. On an active-to-offline
transition, return state variables to a pre-RND configuration. This
example's applications policy does not monitor RND pairs
if it is offline.. */

if (mdi.node_info.service_state == HASIM_NODE_SERVICE_STATE_OFFLINE)
{
  if (mdi.node_info.cache_state != HASIM_NODE_CACHE_INVALID)
  {
    if (mdi.node_info.rg0_state_cache != \
        PMS_PD_NODE_RG0_INFO_GET_STATUS_OFFLINE)
    {
      mdi.node_info.service_state = HASIM_NODE_SERVICE_STATE_ACTIVE;
    }
  }
}
else /* mdi.node_info.service_state == HASIM_NODE_SERVICE_STATE_ACTIVE */
{
  if (mdi.node_info.cache_state != HASIM_NODE_CACHE_INVALID)
  {
    if (mdi.node_info.rg0_state_cache == \
        PMS_PD_NODE_RG0_INFO_GET_STATUS_OFFLINE)
    {

      /* RND */

      for(i=0;i<16;i++)
      {

        if ((mdi.rnd_info[i].receive_state & \
            HASIM_RND_RECEIVE_REGISTERED) != 0)
        {
          pr.session.type = PMS_SR_CALL_NO_RETURN;
          pr.payload.type = PMS_PD_RND_STATUS;
          pr.payload.data.rnd_status.identifier = \
            mdi.mgmt_info.rnd_address_identifier[i];
          pms_receive(&pr, 0, 0);
          mdi.rnd_info[i].receive_state &= !HASIM_RND_RECEIVE_REGISTERED;
        }
      }
    }
  }
}
```

CODE EXAMPLE 7-6 PMS Client Node Interface (Continued)

```
    }

    if ((mdi.rnd_info[i].receive_state & HASIM_RND_MD0_RECEIVE_REGISTERED) \
        != 0)
    {
        pr.session.type = PMS_SR_CALL_NO_RETURN;
        pr.payload.type = PMS_PD_RND_MD0_STATUS;
        pr.payload.data.rnd_status.identifier = \
            mdi.mgmt_info.rnd_address_identifier[i];
        pms_receive(&pr, 0, 0);
        mdi.rnd_info[i].receive_state &= !HASIM_RND_MD0_RECEIVE_REGISTERED;
    }

}

mdi.node_info.service_state = HASIM_NODE_SERVICE_STATE_OFFLINE;
}
}

/* Sync Check */

/* Policy: Sync update checked every SYNCHECK_INTERVAL seconds.. */
if (mdi.node_info.last_sync_check > HASIM_SYNCHECK_INTERVAL)
{
    /* Policy: Don't attempt a sync update if any async partial updates have
       been received within SYNCHECK_INTERVAL.. */
    if (mdi.node_info.last_update > HASIM_SYNCHECK_INTERVAL)
    {
        /* Policy: Don't attempt a sync update if registration for async
           updates have not succeeded.. */
        if (mdi.node_info.receive_state == HASIM_NODE_GROUP_RECEIVE_REGISTERED)
        {
            if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
            {
                if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
                {
                    mdi.node_info.last_sync_check = 0;

                    info_get_fail = 0;

                    /* Get NODE RG0 information.. */
                    ps.session.type = PMS_SR_CALL_RETURN_TIMED;
                    ps.session.info.crt.time = 0;
                    ps.payload.type = PMS_PD_NODE_RG0_INFO_GET_EXECUTE;
```

CODE EXAMPLE 7-6 PMS Client Node Interface (Continued)

```
    if (pms_send(&ps, &pr) == 0)
    {
        if (pr.payload.data.node_rg0_info_get_status.err == \
            PMS_PD_NODE_RG0_INFO_GET_STATUS_SUCCESS)
        {
            rg0_state = pr.payload.data.node_rg0_info_get_status.state;
        }
        else
        {
            info_get_fail = 1;
        }
    }
    else
    {
        info_get_fail = 1;
    }

    /* Get any other NODE info? */

    /* Only mark NODE update as successful if all pieces of data gotten
       were received successfully.. */

    if (info_get_fail == 0)
    {
        mdi.node_info.rg0_state_cache = rg0_state;

        mdi.node_info.last_update = 0;
    }
}
}
}
else
{
    mdi.node_info.last_sync_check = 0;
}
}

/* Validity Check */

/* Process cache state validity transitions. The policy is that on a NODE cache
   transition to invalid, NODE AND RND state variables are returned to an
```


CODE EXAMPLE 7-6 PMS Client Node Interface (Continued)

```
initial configuration.. */

if(mdi.node_info.last_update < HASIM_CHECK_VALID_INTERVAL)
{
    if (mdi.node_info.cache_state != HASIM_NODE_CACHE_VALID)
        mdi.node_info.cache_state = HASIM_NODE_CACHE_VALID;
}
else if((mdi.node_info.last_update >= HASIM_CHECK_VALID_INTERVAL && \
mdi.node_info.last_update < HASIM_CHECK_INVALID_INTERVAL))
{
    if (mdi.node_info.cache_state == HASIM_NODE_CACHE_VALID)
        mdi.node_info.cache_state = HASIM_NODE_CACHE_OLD;
}
else if(mdi.node_info.last_update >= HASIM_CHECK_INVALID_INTERVAL)
{
    if (mdi.node_info.cache_state == HASIM_NODE_CACHE_OLD)
    {

        /* RND */

        for(i=0;i<16;i++)
        {

            if ((mdi.rnd_info[i].receive_state & HASIM_RND_RECEIVE_REGISTERED) != 0)
            {
                pr.session.type = PMS_SR_CALL_NO_RETURN;
                pr.payload.type = PMS_PD_RND_STATUS;
                pr.payload.data.rnd_status.identifier = \
                    mdi.mgmt_info.rnd_address_identifier[i];
                pms_receive(&pr, 0, 0);
                mdi.rnd_info[i].receive_state &= !HASIM_RND_RECEIVE_REGISTERED;
            }

            if ((mdi.rnd_info[i].receive_state & HASIM_RND_MD0_RECEIVE_REGISTERED) \
                != 0)
            {
                pr.session.type = PMS_SR_CALL_NO_RETURN;
                pr.payload.type = PMS_PD_RND_MD0_STATUS;
                pr.payload.data.rnd_status.identifier = \
                    mdi.mgmt_info.rnd_address_identifier[i];
                pms_receive(&pr, 0, 0);
                mdi.rnd_info[i].receive_state &= !HASIM_RND_MD0_RECEIVE_REGISTERED;
            }
        }
    }
}
```

CODE EXAMPLE 7-6 PMS Client Node Interface (Continued)

```
/* NODE*/

if (mdi.node_info.service_state == HASIM_NODE_SERVICE_STATE_ACTIVE)
{
    mdi.node_info.service_state = HASIM_NODE_SERVICE_STATE_OFFLINE;
}

if (mdi.node_info.rg0_app_name_state == \
    HASIM_NODE_RG0_APP_NAME_REGISTERED)
{
    mdi.node_info.rg0_app_name_state = HASIM_NODE_RG0_APP_NAME_UNREGISTERED;
}

if ((mdi.node_info.receive_state & \
    HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED) != 0)
{
    pr.session.type = PMS_SR_CALL_NO_RETURN;
    pr.payload.type = PMS_PD_NODE_RG0_STATUS;
    pms_receive(&pr, 0, 0);
    mdi.node_info.receive_state &= \
        !HASIM_NODE_RG0_STATUS_RECEIVE_REGISTERED;
}

if ((mdi.node_info.receive_state & \
    HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED) != 0)
{
    pr.session.type = PMS_SR_CALL_RETURN_TIMED;
    pr.payload.type = PMS_PD_NODE_RG0_APP_STATE_SET_EXECUTE;
    pms_receive(&pr, 0, 0);
    mdi.node_info.receive_state &= \
        !HASIM_NODE_RG0_APP_STATE_SET_EXECUTE_RECEIVE_REGISTERED;
}

}
}

}
```

The following example shows a PMS client RND interface.

CODE EXAMPLE 7-7 PMS Client RND Interface

```
void
app_hasim_rnd_process(void)
{
    struct pms_send      ps;
    struct pms_receive  pr;

    int                  info_get_fail;

    int                  view;
    int                  md0_config;

    int                  i;

    /* Receive Check */

    /* If RND messages are not receive registered, attempt to register for
    in-use RND address list entries if the service state is active, if
    PMS is in the available state and reachable, and if USER receive messages
    are registered. If registration is successful, force an initial cache
    update.. */

    for(i=0;i<16;i++)
    {
        if (mdi.rnd_info[i].receive_state != HASIM_RND_GROUP_RECEIVE_REGISTERED)
        {
            if (mdi.node_info.service_state == HASIM_NODE_SERVICE_STATE_ACTIVE)
            {
                if (mdi.mgmt_info.rnd_address_identifier[i] != -1)
                {
                    if (mdi.mgmt_info.pms_state == HASIM_MGMT_PMS_STATE_AVAILABLE)
                    {
                        if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
                        {
                            if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
                            {

                                if ((mdi.rnd_info[i].receive_state & \
                                    HASIM_RND_RECEIVE_REGISTERED) == 0)
                                {
                                    pr.session.type = PMS_SR_CALL_NO_RETURN;
                                    pr.payload.type = PMS_PD_RND_STATUS;
                                    pr.payload.data.rnd_status.identifier = \
```

CODE EXAMPLE 7-7 PMS Client RND Interface (Continued)

```
        mdi.mgmt_info.rnd_address_cache[i].identifier;
        if (pms_receive(&pr, app_hasim_receive_post, 0) != -1)
            mdi.rnd_info[i].receive_state |= HASIM_RND_RECEIVE_REGISTERED;
    }

    if ((mdi.rnd_info[i].receive_state & \
        HASIM_RND_MD0_RECEIVE_REGISTERED) == 0)
    {
        pr.session.type = PMS_SR_CALL_NO_RETURN;
        pr.payload.type = PMS_PD_RND_MD0_STATUS;
        pr.payload.data.rnd_md0_status.identifier = \
            mdi.mgmt_info.rnd_address_cache[i].identifier;
        if (pms_receive(&pr, app_hasim_receive_post, 0) != -1)
            mdi.rnd_info[i].receive_state |= \
                HASIM_RND_MD0_RECEIVE_REGISTERED;
    }

    /* Force an info_get immediately after registering.. */
    mdi.rnd_info[i].last_sync_check = HASIM_SYNCHECK_INTERVAL;
}
}
}
}
}

/* Sync Check */

for(i=0;i<16;i++)
{
    /* Policy: Sync update checked every SYNCHECK_INTERVAL seconds.. */
    if (mdi.rnd_info[i].last_sync_check > HASIM_SYNCHECK_INTERVAL)
    {
        /* Policy: Don't attempt a sync update if any async partial updates have
           been received within SYNCHECK_INTERVAL.. */
        if (mdi.rnd_info[i].last_update > HASIM_SYNCHECK_INTERVAL)
        {
            /* Policy: Don't attempt a sync update if registration for async
               updates have not succeeded.. */
            if (mdi.rnd_info[i].receive_state == HASIM_RND_GROUP_RECEIVE_REGISTERED)
            {
                if (mdi.user_info.receive_state == HASIM_USER_RECEIVE_REGISTERED)
                {
                    if (mdi.user_info.pms_view == HASIM_USER_PMS_VIEW_REACHABLE)
```

CODE EXAMPLE 7-7 PMS Client RND Interface (*Continued*)

```
{
mdi.rnd_info[i].last_sync_check = 0;

info_get_fail = 0;

/* Get RND information.. */
ps.session.type = PMS_SR_CALL_RETURN_TIMED;
ps.session.info.crt.time = 0;
ps.payload.type = PMS_PD_RND_INFO_GET_EXECUTE;
ps.payload.data.rnd_info_get_execute.identifier = \
    mdi.mgmt_info.rnd_address_identifier[i];

if (pms_send(&ps, &pr) == 0)
{
    if (pr.payload.data.rnd_info_get_status.err == \
        PMS_PD_RND_INFO_GET_STATUS_ERR_NONE)
    {
        view = pr.payload.data.rnd_info_get_status.view;
    }
    else
    {
        info_get_fail = 1;
    }
}
else
{
    info_get_fail = 1;
}

/* Get RND MD0 information.. */
ps.session.type = PMS_SR_CALL_RETURN_TIMED;
ps.session.info.crt.time = 0;
ps.payload.type = PMS_PD_RND_MD0_INFO_GET_EXECUTE;
ps.payload.data.rnd_md0_info_get_execute.identifier = \
    mdi.mgmt_info.rnd_address_identifier[i];

if (pms_send(&ps, &pr) == 0)
{
    if (pr.payload.data.rnd_md0_info_get_status.err == \
        PMS_PD_RND_MD0_INFO_GET_STATUS_ERR_NONE)
    {
        md0_config = pr.payload.data.rnd_md0_info_get_status.config;
    }
    else
    {
        info_get_fail = 1;
    }
}
```

CODE EXAMPLE 7-7 PMS Client RND Interface (*Continued*)

```
    }
    else
    {
        info_get_fail = 1;
    }

    /* Only mark MGMT update as successful if all pieces of data
       were received successfully.. */

    if (info_get_fail == 0)
    {
        mdi.rnd_info[i].view_cache = view;

        mdi.rnd_info[i].md0_config_cache = md0_config;

        mdi.rnd_info[i].last_update = 0;
    }
}
}
}
else
{
    mdi.rnd_info[i].last_sync_check = 0;
}
}

/* Validity Check */

/* Process cache state validity transitions. The policy is on a RND cache
   transition to invalid, return RND state variables for the pair to an initial
   configuration.. */

for(i=0;i<16;i++)
{
    if(mdi.rnd_info[i].last_update < HASIM_CHECK_VALID_INTERVAL)
    {
        if (mdi.rnd_info[i].cache_state != HASIM_RND_CACHE_VALID)
            mdi.rnd_info[i].cache_state = HASIM_RND_CACHE_VALID;
    }
    else if((mdi.rnd_info[i].last_update >= HASIM_CHECK_VALID_INTERVAL && \
            mdi.rnd_info[i].last_update < HASIM_CHECK_INVALID_INTERVAL))
    {
```

CODE EXAMPLE 7-7 PMS Client RND Interface (*Continued*)

```
    if (mdi.rnd_info[i].cache_state == HASIM_RND_CACHE_VALID)
        mdi.rnd_info[i].cache_state = HASIM_RND_CACHE_OLD;
    }
else if(mdi.rnd_info[i].last_update >= HASIM_CHECK_INVALID_INTERVAL)
    {
    if (mdi.rnd_info[i].cache_state == HASIM_RND_CACHE_OLD)
        {

        /* RND */

        if ((mdi.rnd_info[i].receive_state & HASIM_RND_RECEIVE_REGISTERED) != 0)
            {
            pr.session.type = PMS_SR_CALL_NO_RETURN;
            pr.payload.type = PMS_PD_RND_STATUS;
            pr.payload.data.rnd_status.identifier = \
                mdi.mgmt_info.rnd_address_identifier[i];
            pms_receive(&pr, 0, 0);
            mdi.rnd_info[i].receive_state &= !HASIM_RND_RECEIVE_REGISTERED;
            }

            if ((mdi.rnd_info[i].receive_state & \
                HASIM_RND_MD0_RECEIVE_REGISTERED) != 0)
                {
                pr.session.type = PMS_SR_CALL_NO_RETURN;
                pr.payload.type = PMS_PD_RND_MD0_STATUS;
                pr.payload.data.rnd_status.identifier = \
                    mdi.mgmt_info.rnd_address_identifier[i];
                pms_receive(&pr, 0, 0);
                mdi.rnd_info[i].receive_state &= !HASIM_RND_MD0_RECEIVE_REGISTERED;
                }

            }
        }
    }
}
```


Solaris Operating System APIs

This chapter introduces Solaris Operating System (Solaris OS) APIs of concern to the Netra CT server, including configuration and status of the system frutree and environmental monitoring with sensor status information. This is handled through the Platform Information and Control Library (PICL) framework, gathering FRU-ID information, and dynamic reconfiguration interfaces. These subjects are addressed in:

- [“Solaris Operating System PICL Framework” on page 138](#)
- [“PICL Frutree Topology” on page 140](#)
- [“PICL Man Page References” on page 146](#)
- [“Dynamic Reconfiguration Interfaces” on page 148](#)
- [“Programming Temperature Sensors Using the PICL API” on page 151](#)
- [“Programming Watchdog Timers Using the PICL API” on page 153](#)
- [“Displaying FRU-ID Data” on page 156](#)
- [“MCNet Support” on page 159](#)

Solaris Operating System PICL Framework

PICL provides a method to publish platform-specific information for clients to access in a way that is not specific to the platform. The Solaris PICL framework provides information about the system configuration which it maintains in the PICL tree. Within this PICL tree is a subtree named *frutree*, that represents the hierarchy of system FRUs with respect to a root node in the tree called *chassis*. The *frutree* represents physical resources of the system.

The main components of the PICL framework are:

- PICL interface (`libpicl.so`) – Implements the generic platform-independent interface that clients can use to access the platform-specific information.
- PICL tree (`libpicltree.so`) – A repository of all the nodes and properties representing the platform configuration.
- PICL plug-in modules – Shared objects that publish platform-specific data in the PICL tree.
- PICL daemon (`picld`) – Maintains and controls access to the PICL information from clients and from PICL plug-in modules.

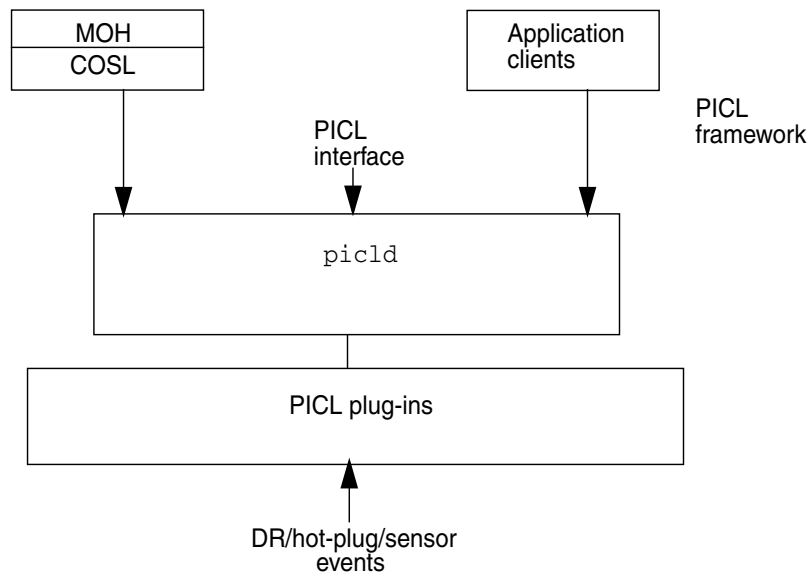


FIGURE 8-1 PICL Daemon (`picld`) and Plug-Ins

FIGURE 8-1 diagrams the PICL daemon (`picld`) and its connection to the PICL plug-ins, some of which are common to all platforms. These plug-ins are responsible for populating and maintaining the PICL tree during system boot and dynamic reconfiguration (DR) operations. They also handle sensor events.

Application clients use `libpicl(3LIB)` to interface with `picld` to display or set parameters using Managed Object Hierarchy (MOH), RMI, SNMP, or Solaris PICL client interfaces. MOH uses the common operating system library (COSL) interface to access `picld`, which in turn uses the `libpicl` interfaces.

Updates to the system frutree are done during DR operations, which are performed using `cfgadm(1M)`, or during `fru latch` and `unlatch` operations.

The following section identifies the exported interfaces in the PICL tree which are basically the nodes and the properties of the node that are present in the PICL tree.

PICL Frutree Topology

To read the PICL frutree data of the system, use the `prtpicl(1M)` command. The structure of the PICL frutree involves a hierarchical representation of nodes. The immediate descendants of `/frutree` are one or more `fru` nodes by the name of chassis.

FRUs and their locations are represented by nodes of classes `fru` and `location` respectively, in the PICL frutree under the `/frutree` node. The `port` node is an extension of the `fru` class.

The three major node classes, `location`, `fru`, and `port`, are summarized in [TABLE 8-1](#). Each of these classes is populated with various properties, among which are `State` and `Condition`. More detailed information is provided in the sections following this summary table.

TABLE 8-1 PICL FRUtree Topology Summary

Node Class	Properties	Description
location	SlotType	Type of location.
	Label	Slot/Label information.
	GeoAddr	Geographical address.
	StatusTime	Time when <code>State</code> was updated last.
	Bus-addr	Bus address.
	State	State of the location: <code>empty</code> , <code>connected</code> , <code>disconnected</code> , or <code>unknown</code> .
fru	FruType	Type of FRU.
	Devices	Table of node handles in platform tree.
	State	State of the FRU – <code>configured</code> , <code>unconfigured</code> , or <code>unknown</code> .
	StatusTime	Time when <code>State</code> was updated last.
	Condition	Condition or operational state of the FRU – <code>ok</code> , <code>failing</code> , <code>failed</code> , <code>unknown</code> , or <code>unusable</code> .
	ConditionTime	Time when <code>Condition</code> was updated last.
port	Bus-addr	Bus address of port – <code>network</code> , <code>serial</code> , or <code>parallel</code> .
	GeoAddr	Geographical address of port.
	Label	Label information.
	PortType	Type of port.

TABLE 8-1 PICL FRUtree Topology Summary

Node Class	Properties	Description
	State	State of the port – up, down, or unknown.
	StatusTime	Time when State was updated last.
	Condition	Condition of the port – ok, unknown, failing, or failed.
	ConditionTime	Time when Condition was updated last.
	Devices	Table of node handles in platform device tree.

Chassis Node Property Updates

In addition to those properties already defined by PICL, the following property is added:

ChassisType

CHARSTRING read-only

The `ChassisType` read-only property represents the type of chassis. The value of `ChassisType` is currently defined as: `SUNW,NetraCT-410` for a 5-slot chassis or `SUNW,NetraCT-810` for an 8-slot chassis.

The value of this property can be derived using the following:

```
prtpicl -v -c fru | grep ChassisType | cut -f2 | tr -d ' '
```

There should be a configuration file of this name with the `.conf` extension in the `/usr/platform/'uname -i'/lib/picl/plugins/` directory. If none is provided, then the `frutree` is not initialized.

fru Class Properties

Where the following `fru` class properties are writeable, permission checks govern that they be written to by a process with the user ID of root.

fru State

CHARSTRING read-only

The `State` property of the `fru` class node represents the occupant state of the `cfgadm` attachment point associated with `fru` node. In such a case, a read operation of this property directs the plug-in to check the state of the occupant using `libcfgadm` to determine the latest `State` information.

The various state values are shown in [TABLE 8-2](#).

TABLE 8-2 PICL FRU State Value Properties

State	Property Value Description
<code>unconfigured</code>	The FRU is not configured and unusable. See <code>cfgadm(1M)</code> for details.
<code>configured</code>	The FRU is configured and usable. See <code>cfgadm(1M)</code> for details.

FRU Condition

CHARSTRING read-only

The `Condition` property of the `fru` class node represents the condition of occupant of the `cfgadm` attachment point. The various condition values are shown in [TABLE 8-3](#). When `libcfgadm` interfaces are not available, a platform must provide the same semantics using platform-specific interfaces in defining this property.

TABLE 8-3 PICL FRU Condition Value Properties

Condition	Property Value Description
<code>unknown</code>	FRU condition could not be determined. See <code>cfgadm(1M)</code> for details.
<code>ok</code>	FRU is functioning as expected. See <code>cfgadm(1M)</code> for details.
<code>failing</code>	A recoverable fault was found. See <code>cfgadm(1M)</code> for details.
<code>failed</code>	An unrecoverable fault was found. See <code>cfgadm(1M)</code> for details.
<code>unusable</code>	FRU is unusable for undetermined reason. See <code>cfgadm(1M)</code> for details.

The FRU Condition can represent both software and hardware faults.

port Class Node

The connectivity between nodes in a telecommunications network is established by a link that provides the physical transmission medium. A `port` node represents a resource of a `fru` that provides such a link. Examples of ports are `serial` port and `network` port.

The `port` class node extends the PICL frutree definition of `fru` class of nodes. A `port` is always a child of a `fru` class, even if it is the only resource of the `fru`. There are no location or `fru` nodes beneath a `port` class node, because FRUs linked to the `port` class node are not managed in the domain in which `port` class node exists. There might be dependencies, such as when a remote device is cabled to a `port` node. These dependencies can influence the state of the `port`, but not necessarily the FRU itself.

The PICL frutree plug-in is responsible for identifying the `port` class nodes and creating the respective nodes in the frutree.

Note – The `port` class node should not be associated with USB port or SCSI port. These are locations into which a FRU can be plugged, become visible to the system CPU, and managed by it. FRUs beyond the `port` class of nodes are not visible to the CPU.

port Class Properties

`port` class properties consist of `State` and `Condition`, values of which are shown in the following paragraphs.

State

CHARSTRING read-only

A `port` class node can be in one of the states shown in [TABLE 8-4](#):

TABLE 8-4 Port Class State Values

Port State Values	Description
down	A port is down when its link state is down, that is, a carrier was not detected.
up	A port is up when its link state is up, that is, a carrier is detected.
unknown	The plug-in cannot determine the state of the port.

The state of the `port` node is maintained by the frutree plug-in. The `State` value is initially determined by looking at the `kstat` information published by the device driver that owns the port. If the device driver information is not determined, this value remains unknown. The parent `fru` of the `port` must set its state to configured for the `port` to be anything other than unknown. See `kstat(1M)` for details.

Condition

CHARSTRING read-write

The `Condition` value of a `port` class node carries the same meaning as the `cfgadm` value of the attachment point, as shown in [TABLE 8-5](#).

TABLE 8-5 Port Condition Values

Port Condition Values	Description
<code>ok</code>	Port is functioning as expected.
<code>failing</code>	A predictive failure has been detected. This typically occurs when the number of correctable errors exceeds a threshold.
<code>failed</code>	Port has failed. It cannot transmit or receive data due to an internal fault. This indicates a broken path within the FRU, and not external to the FRU which would be denoted by its link state.
<code>unknown</code>	Port condition could not be determined.

Initial `Condition` values can be obtained by looking at the driver `kstat` information, if present. A device driver managing a resource of the FRU can influence the overall condition of the FRU by sending appropriate fault events. The property information is valid only when the parent `fru` state is configured.

PortType

CHARSTRING read-only

This `PortType` property indicates the functional class of `port` device, as shown in [TABLE 8-6](#).

TABLE 8-6 PortType Property Values

PortType Values	Description
<code>network</code>	Represents a network device.
<code>serial</code>	Represents a serial device.
<code>parallel</code>	Represents a parallel port.

Common Property Updates

The following properties are common to all PICL classes:

GeoAddr

UINT read-only

This property indicates the geographical address of the node in relation to its parent node. It should be possible to point to the physical location (slot number) of the node in its parent domain. For example, the Netra CT 810 server describes a location's `GeoAddr` under the chassis node as its physical slot number. This could differ from the `Label` information printed on the chassis itself. In this instance, the system controller slot on the Netra CT 810 system chassis is labelled as `CPU`, although its `GeoAddr` has a value of 1. Note that the `Label` property might not have the physical slot number embedded in it.

StatusTime

TIMESTAMP read-only

This property indicates when the `State` property was last updated. This can indicate when a FRU was last inserted or removed, configured or unconfigured, or when the port link went down. Status time is updated even for transitional state changes.

ConditionTime

TIMESTAMP read-only

This property indicates when the `Condition` property was last updated. Using this property, for example, system management software can calculate how long a `fru` or `port` has been in operation before failure.

Temperature Sensor Node State

CHARSTRING

A temperature sensor node is in the PICL frutree under the `Environment` property of the `fru` node.

The temperature sensors are represented as `PICL_CLASS_TEMPERATURE_SENSOR` class in the PICL tree. A `State` property is declared for each temperature sensor node representing the state information as shown in [TABLE 8-7](#).

TABLE 8-7 State Property Values for Temperature Sensor Node

State Property Values	Description
ok	Environment state is OK.
warning	Environment state is warning, (that is, current temperature is below lower or above upper warning temperature).
failed	Environment state is failed (that is, current temperature is below lower or above upper critical temperature).
unknown	Environment state is unknown (that is, current temperature cannot be determined).

PICL Man Page References

[TABLE 8-8](#) lists the Solaris OS man pages that document the PICL framework and API. You can view the following man pages at the command line or on the Solaris OS documentation web site (<http://docs.sun.com/documentation>).

TABLE 8-8 PICL Man Pages

Man Page	Description
<code>picld(1M)</code>	Describes how the daemon initializes plug-in modules at startup. The man page also describes the PICL tree and PICL plug-in modules.
<code>libpicl(3LIB)</code>	Lists the library functions clients use to interface with the PICL daemon in order to access information from the PICL tree.
<code>libpicl(3PICL)</code>	Client API for sending requests to the PICL daemon to access the PICL tree.

TABLE 8-8 PICL Man Pages (*Continued*)

Man Page	Description
<code>picld_log(3PICLTREE)</code>	Describes the function the PICL daemon and the plug-in modules use to log messages and inform users of any error or warning conditions.
<code>picl_plugin_register(3PICLTREE)</code>	Describes the function plug-in modules use to register itself with the PICL daemon.
<code>prtpicl(1M)</code>	Prints the PICL tree. The <code>prtpicl</code> command prints the PICL tree maintained by the PICL daemon. The output of <code>prtpicl</code> includes the name and PICL class of the nodes.
(3LIB) Functions	
<code>picl_initialize(3PICL)</code>	Initiates a session with the PICL daemon.
<code>picl_get_first_prop(3PICL)</code>	Gets a property handle of a node.
<code>picl_get_next_by_col(3PICL)</code>	Accesses a table property.
<code>picl_get_next_by_row(3PICL)</code>	Accesses a table property.
<code>picl_get_next_prop(3PICL)</code>	Gets a property handle of a node.
<code>picl_get_prop_by_name(3PICL)</code>	Gets the handle of the property by name.
<code>picl_get_propinfo(3PICL)</code>	Gets the information about a property.
<code>picl_get_propinfo_by_name(3PICL)</code>	Gets property information and handle of a property by name.
<code>picl_get_propval(3PICL)</code>	Gets the value of a property.
<code>picl_get_propval_by_name(3PICL)</code>	Gets the value of a property by name.
<code>picl_get_root(3PICL)</code>	Gets the root handle of the PICL tree.
<code>picl_set_propval(3PICL)</code>	Sets the value of a property to the specified value.
<code>picl_set_propval_by_name(3PICL)</code>	Sets the value of a named property to the specified value.
<code>picl_shutdown(3PICL)</code>	Shuts down the session with the PICL daemon.
<code>picl_strerror(3PICL)</code>	Gets error message string.
<code>picl_wait(3PICL)</code>	Waits for PICL tree to refresh.
<code>picl_walk_tree_by_class(3PICL)</code>	Walks subtree by class.

For examples of use of these functions, see [“Programming Watchdog Timers Using the PICL API” on page 153](#).

Dynamic Reconfiguration Interfaces

The dynamic reconfiguration (DR) interfaces allow resources to be reconfigured without user intervention when system resources are added or removed while the system is running. Traditionally, applications assume that OS resources remain static after boot. In DR situations, challenges faced by applications include the following:

- Addition or availability of new devices. Applications might want to be notified in order to make use of the newly added resources.
- Removal of devices. Applications need to be notified of pending resource removal from the system so they can either block or prepare for the pending operation.

The Solaris OS has knowledge of DR operations, but certain applications might not. If an application is holding the resources involved in the DR operation, the operation will fail. To be successful, applications need to be dynamically aware of the current state of the system. The Solaris DR framework includes the Reconfiguration Coordination Manager (RCM), `cfgadm(1m)`, and `libcfgadm(3LIB)`. It also includes the PCI hot-plug/cPCI Hot-swap framework (`cfgadm_pci(1M)`), SCSI hot-plug framework (`cfgadm_scsi(1M)`), and the Hot-swap Controller driver (`cphsc(7D)`).

The following sections describe:

- [“Reconfiguration Coordination Manager” on page 148](#)
- [“Hot-Swap Support” on page 149](#)
- [“Configuration Administration” on page 150](#)

Reconfiguration Coordination Manager

The Reconfiguration Coordination Manager (RCM) is a generic framework which allows DR to interact with system management software. The framework enables automated DR removal operations on platforms with proper software and hardware configuration. RCM defines generic APIs to coordinate DR operations between DR initiators and DR clients during resource removal. For details on RCM, go to <http://www.sun.com/documentation>.

Hot-Swap Support

The Netra CT server supports the following three hot-swap models according to the PICMG CompactPCI Hot-Swap specifications version 2.1 R1.0:

- Basic hot-swap
- Full hot-swap
- High availability hot swap

These models can be described by two terms:

- Hardware connection process – the electrical connection (and disconnection) of an I/O board
- Software connection process – the software configuration (and unconfiguration) of the I/O board by the operating system (allocating or releasing PCI resources, attaching or detaching device drivers, and so on.)

In the basic hot-swap model, the hardware connection process can be performed automatically by the hardware, while the software connection process requires operator assistance.

In the full hot-swap model, both the hardware and the software connection processes are performed automatically. The Netra CT server is configured for full hot swap by default. The mode of a slot can be reconfigured to basic hot swap using the `cfgadm` command in cases where a third-party board does not support full hot swap.

In the high-availability model, software has the capability of controlling the power-on of the FRU hardware, beyond the hardware and software connection processes. Drivers and services can isolate a board from the system until an operator is able to intervene.

The Netra CT server uses the `cfgadm(1M)` utility for administering the hot-swap process. This includes connecting and disconnecting, configuring, and unconfiguring the hardware and software, and setting various operation modes. Elements of the `cfgadm(1M)` utility are described in the next section.

On the Netra CT server, CPU card, CPU transition card, and I/O board hot swapping is supported. It should be noted that non-hotswap friendly devices can be supported only in basic hot-swap mode. See the *Netra CT Server Service Manual* (819-2743) for list of hot-swappable FRUs.

Configuration changes are handled in a coherent way, because DR and the Frutree management framework are integrated in PICL. PICL frutree properties and `cfgadm` attachment point elements are mapped one-to-one, which creates data consistency. All DR operations are coordinated with a service processor.

Configuration Administration

Configuration administration of a dynamically reconfigurable system is carried out through `cfgadm(1M)`, which can display status, invoke configuration state changes, and invoke hardware specific functions. See the *Netra CT System Administration Guide* (819-2743) for more information on the `cfgadm` utility.

The `libcfgadm(3LIB)` command can be used to display a library of configuration interfaces.

Use `cfgadm` to perform a connect operation on a cPCI FRU, for example:

- To power on a FRU
- To check for HEALTHY#
- To bring a FRU out of reset and connect it electrically to the cPCI backplane

```
# cfgadm -c connect operation
```

Use `cfgadm` to perform a disconnect operation on a cPCI FRU, for example:

- To notify applications (via RCM)
- To bring the FRU bridge in reset and disconnecting it electrically from the cPCI backplane
- To power OFF a FRU

```
# cfgadm -c disconnect operation
```

Programming Temperature Sensors Using the PICL API

Temperature sensor states can be read using the `libpicl` API. The properties that are supported in a PICL temperature sensor class node are listed in [TABLE 8-9](#).

TABLE 8-9 PICL Temperature Sensor Class Node Properties

Property	Type	Description
<code>LowWarningThreshold</code>	INT	Low threshold for warning
<code>LowShutdownThreshold</code>	INT	Low threshold for shutdown
<code>LowPowerOffThreshold</code>	INT	Low threshold for power off
<code>HighWarningThreshold</code>	INT	High threshold for warning
<code>HighShutdownThreshold</code>	INT	High threshold for shutdown
<code>HighPowerOffThreshold</code>	INT	High threshold for power off

The PICL plug-in receives these sensor events and updates the `State` property based on the information extracted from the IPMI message. It then posts a PICL event.

The threshold levels of the PICL node class `temperature-sensor` are:

- `Warning`
- `Shutdown`
- `PowerOff`

[TABLE 8-10](#) lists the PICL threshold levels and their MOH equivalents.

TABLE 8-10 PICL Threshold Levels and MOH Equivalents

PICL Threshold levels	MOH Equivalent
<code>LowWarningThreshold</code>	<code>LowerThresholdNonCritical</code>
<code>LowShutdownThreshold</code>	<code>LowerThresholdCritical</code>
<code>LowPowerOffThreshold</code>	<code>LowerThresholdFatal</code>
<code>HighWarningThreshold</code>	<code>UpperThresholdNonCritical</code>
<code>HighShutdownThreshold</code>	<code>UpperThresholdCritical</code>
<code>HighPowerOffThreshold</code>	<code>UpperThresholdFatal</code>

To obtain a reading of temperature sensor states, type the `prtpicl -v` command:

```
# prtpicl -c temperature-sensor -v
```

PICL output of the temperature sensors on a Netra CT system using a Netra-CP2500 host is shown in [CODE EXAMPLE 8-1](#).

CODE EXAMPLE 8-1 Example Output of PICL Temperature Sensors

```
# prtpicl -c temperature-sensor -v
CPU-sensor (temperature-sensor, 41000003fb)
:Condition      ok
:HighPowerOffThreshold      115
:HighShutdownThreshold     110
:HighWarningThreshold      105
:LowPowerOffThreshold       -20
:LowShutdownThreshold       -10
:LowWarningThreshold        -5
:Temperature      78
:Label            Ambient
:GeoAddr          0xe
:_class           temperature-sensor
:name             CPU-sensor
```

Note – PICL clients can use the `libpicl` APIs to set and get various properties of this sensor.

Programming Watchdog Timers Using the PICL API

The Netra CT system's watchdog service captures catastrophic faults in the Solaris OS running on either a host or satellite CPU board. The watchdog service reports such faults to the alarm card by means of either an IPMI message or by a de-assertion of the CPU's HEALTHY# signal.

The Netra CT system management controller provides two watchdog timers, the watchdog level 2 (WD2) timer and the watchdog level 1 (WD1) timer. Systems management software starts and the Solaris OS periodically pats the timers before they expire. If the WD2 timer expires, the watchdog function of the WD2 timer forces the SPARC[®] processor to optionally reset. The maximum range for WD2 is 255 seconds.

The WD1 timer is typically set to a shorter interval than the WD2 timer. User applications can examine the expiration status of the WD1 timer to get advance warning if the main timer, WD2, is about to expire. The system management software has to start WD1 before it can start WD2. If WD1 expires, then WD2 starts only if enabled. The maximum range for WD1 is 6553.5 seconds.

The watchdog subsystem is managed by a PICL plug-in module. This PICL plug-in module provides a set of PICL properties to the system, which enables a Solaris PICL client to specify the attributes of the watchdog system.

To use the PICL API to set the watchdog properties, your application must adhere to the following sequence:

1. Before setting the watchdog timer, use the PMS API to disable the primary HEALTHY# signal monitoring for the CPU board on which the watchdog timer is to be changed.

To do this, switch to the alarm card CLI and use the command `pmsd infoshow`, specifying the slot number. The output will indicate whether the card is in MAINTENANCE mode or OPERATIONAL mode.

```
# pmsd infoshow -s slot_number
config=<MAINTENANCE|OPERATIONAL>
ALARM_STATE=NONE
```

If the card is in OPERATIONAL mode, switch it into MAINTENANCE mode by issuing the following command:

```
# pmsd operset -s slot_number -o MAINT_CONFIG
```

This disables the primary HEALTHY# signal monitoring of the board in the specified slot.

2. In your application, use the PICL API to disarm, set, and arm the active watchdog timer.

Refer to the `picld(1M)`, `libpicl(3LIB)`, and `libpicl(3PICL)` man pages for a complete description of the PICL architecture and programming interface. Develop your application using the PICL programming interface to do the following:

- Disarm the active watchdog timer.
- Change the watchdog timer PICL properties to the required values.
- Re-arm the watchdog timer. The properties of `watchdog-controller` and `watchdog-timer` are defined in [TABLE 8-11](#), [TABLE 8-12](#), and [TABLE 8-13](#).

3. Use the PMS API to enable the primary HEALTHY# signal monitoring on the CPU card in the specified slot.

From the alarm card CLI, switch the card back to OPERATIONAL mode by issuing the following command:

```
# pmsd operset -s slot_number -o OPER_CONFIG
```

HEALTHY# monitoring will be enabled again on the card in the slot that you specified.

Refer to [Chapter 7](#) for information on Processor Management Services (PMS).

PICL interfaces for the watchdog plug-in module (see [TABLE 8-11](#)) include the nodes `watchdog-controller` and `watchdog-timer`.

TABLE 8-11 Watchdog Plug-in Interfaces for Netra CT 810 and 410 Server Software

PICL Class	Property	Meaning
<code>watchdog-controller</code>	<code>WdOp</code>	Represents a watchdog subsystem.
<code>watchdog-timer</code>	<code>State</code>	Represents a watchdog timer hardware that belongs to its controller. Each timer depends on the status of its peers to be activated or deactivated.
	<code>WdTimeout</code>	Timeout for the watchdog timer.
	<code>WdAction</code>	Action to be taken after the watchdog expires.

TABLE 8-12 Properties Under `watchdog-controller` Node

Property	Operations	Description
<code>WdOp</code>	<code>arm</code>	Activates all timers under the controller with values already set for <code>WdTimeout</code> and <code>WdAction</code> .
	<code>disarm</code>	All active timers under the controller will be stopped.

TABLE 8-13 Properties Under `watchdog-timer` Node

Property	Values	Description
<code>State</code>	<code>armed</code>	Indicates timer is armed or running. Cleared by <code>disarm</code> .
	<code>expired</code>	Indicates timer has expired. Cleared by <code>disarm</code> .
	<code>disarmed</code>	Default value set at boot time. Indicates timer is disarmed or stopped.
<code>WdTimeout</code> *	Varies by system and timer level	Indicates the timer initial countdown value. Should be set prior to arming the timer.
<code>WdAction</code> †	<code>none</code>	Default value. No action is taken.
	<code>alarm</code>	Send notifications to system alarm hardware by means of <code>HEALTHY#</code> .
	<code>reset</code>	Perform a soft or hard reset the system (implementation specific).
	<code>reboot</code>	Reboot the system.

* A platform might not support a specified timeout resolution. For example Netra CT systems only take -1, 0, and 100 6553500 ms in increments of 100 msec. (Level 1), and -1 -255 seconds (Level 2).

† A specific timer node might not support all action types. For example Netra CT watchdog level 1 timer supports only “none” and “alarm” actions. Watchdog level 2 timer supports only “none” and “reset”

To identify current settings of `watchdog-timer`, issue the command `prtpicl -v` as shown in [CODE EXAMPLE 8-2](#).

CODE EXAMPLE 8-2 Example of `watchdog-timer`

```
# prtpicl -v -c watchdog-timer
watchdog-level1 (watchdog-timer, 370000058e)
:WdAction      alarm
:WdTimeout     0x2710
:State         armed
:_class        watchdog-timer
:name          watchdog-level1
watchdog-level2 (watchdog-timer, 3700000591)
:WdAction      none
:WdTimeout     0xffffffff
:State         disarmed
:_class        watchdog-timer
:name          watchdog-level2
```

Displaying FRU-ID Data

Sun FRU-ID is the container for presenting the FRU-ID data. If the Sun FRU-ID container is not present, the FRU-ID Access plug-in looks for the IPMI FRU-ID container of cPCI FRUs. It then converts FRU-ID data from IPMI format to Sun FRU-ID format and presents the result in Sun FRU-ID ManR (manufacturer record) format.

The command `prtfriu(1M)` displays FRU data of all FRUs in the PICL fru tree. When `prtfriu` is run on the host CPU, FRU data is displayed for the host CPU and the satellite CPUs.) [CODE EXAMPLE 8-3](#) shows an example of the output of the `prtfriu` command.

CODE EXAMPLE 8-3 Sample Output of `prtfriu` Command

```
# prtfriu
/frutree
/frutree/chassis (fru)
/frutree/chassis/AL-1?Label=AL 1
/frutree/chassis/AL-1?Label=AL 1/AL-1 (fru)
/frutree/chassis/IO-2?Label=I.O 2
```


CODE EXAMPLE 8-3 Sample Output of prtfru Command (Continued)

```
/ECO_CurrentR/Hardware_Revision: 00
/ECO_CurrentR/HW_Dash_Level: 00
/CPUFirmwareR
/CPUFirmwareR/UNIX_Timestamp32: Thu Sep 15 18:02:32 PDT 2005
/CPUFirmwareR/CPU_FW_Part_No: 5252225
/CPUFirmwareR/CPU_FW_Dash_Level: 01
/Drawer_InfoR
/Drawer_InfoR/UNIX_Timestamp32: Thu Sep 15 18:02:32 PDT 2005
/Drawer_InfoR/Drawer_Id: 000000
/Drawer_InfoR/Drawer_Type: 0000000000000000
/Drawer_InfoR/Access_Model: 0000000000000000
/Drawer_InfoR/Slot_Mode: 0000000000000000
/Drawer_InfoR/Reserved_Data:
00000000000000000000000000000000000000000000000000000000000000
00000000000000000000000000000000000000000000000000000000000000
/Customer_DataR
/Customer_DataR/UNIX_Timestamp32: Thu Sep 15 18:02:32 PDT
2005
/Customer_DataR/Cust_Data: Customer Data
SEGMENT: SD
/ManR
/ManR/UNIX_Timestamp32: Thu Sep 15 18:02:32 PDT 2005
/ManR/Fru_Description: FRUID,PRGM INSTR,MBD,JADE
/ManR/Manufacture_Loc: CELESTICA,THAILAND
/ManR/Sun_Part_No: 5017031
/ManR/Sun_Serial_No: JM0001
/ManR/Vendor_Name: Celestica
/ManR/Initial_HW_Dash_Level: 06
/ManR/Initial_HW_Rev_Level: 06
/ManR/Fru_Shortname: CPU
/SpecPartNo: 885-0530-01
/frutree/chassis/RTM?Label=RTM
/frutree/chassis/RTM?Label=RTM/RTM (container)
SEGMENT: SD
/ManR
/ManR/UNIX_Timestamp32: Tue Aug 23 18:51:25 PDT 2005
/ManR/Fru_Description: FRUID,PRGM INSTR,RTM,JADE_HOST
/ManR/Manufacture_Loc: CELESTICA,THAILAND
/ManR/Sun_Part_No: 5017032
/ManR/Sun_Serial_No: 000001
/ManR/Vendor_Name: Celestica
/ManR/Initial_HW_Dash_Level: 01
/ManR/Initial_HW_Rev_Level: 01
/ManR/Fru_Shortname: RTM-H
/SpecPartNo: 885-0531-01
/frutree/chassis/c0::dsk.c0t0d0?Label=HDD 0
```

CODE EXAMPLE 8-3 Sample Output of `prtfru` Command (Continued)

```
/frutree/chassis/c0::dsk.c0t0d0?Label=HDD 0/c0::dsk.c0t0d0 (fru)
rwings-110#
```

MCNet Support

Communication between CPUs is enabled by MCNet (`mcn(7D)`), which presents an Ethernet-like interface over the cPCI bus in accordance with PICMG 2.14. The interface is configured automatically during system boot, and supports all existing network tools, such as `ifconfig(1M)`, `netstat(1M)` and so forth. The CPUs must be MCNet-capable in order to communicate with each other.

Glossary

A

- AC** Alarm card. The alarm card is used in the Netra 810 and Netra CT 410 servers to provide system control functions. The alarm card resides in slot 8 in the Netra CT 810 server and in slot 1 in the Netra 410 server.
- ACL** Access control list; a file that details which SNMP management applications can access information maintained by the MOH. The file also lists which hosts can receive SNMP traps or events.
- alarm severity profile** A managed entity that contains the severity assignments for the reported alarms.
- ASN1** Abstract notation one. The notation used in a text file for a MIB. The variables containing the information that SNMP can access are described in this file.
- attribute value change record** A managed entity used to represent logged information resulting from attribute value change notifications. Instances of this managed entity are created automatically by the network entity (NE), and deleted by the NE or by request of the managing system.

C

- CGTP** Carrier Grade Transport Protocol. CGTP network interfaces send and receive packets on redundant networks. These software devices use CGTP protocol. See the `ifcgtp(7)` man page, which details the general properties of the network interfaces.
- CLI** Command-line interface. The primary user interface to the alarm card.
- cPCI** Compact PCI.

E

- EFDMBean** Event Forwarding Discriminator. A managed entity used as a notification forwarder discriminator. At startup it registers itself as a listener to all the broadcaster MBeans registered with the MBeanServer, then listens for MBeanServer creation notifications to register with newly created MBeans.
- equipment** A managed entity used to represent the various externally manageable physical components of the network entity (NE) that are not modeled using the Plug-in Unit or Equipment Holder managed entities.
- equipment holder** A managed entity representing physical resources of the NE that are capable of holding other physical resources. An instance of this managed entity exists for each rack, shelf, drawer, and slot of the NE.

F

- full log** A managed entity used to group multiple instances of the Managed Entity Creation Log Record, Managed Entity Deletion Log Record, State Change Log Record, Attribute Value Change Log Record, and/or Alarm Record managed entities to form a log. This managed entity contains information that, among other things, allows the management system to control the behavior of the log.

G

GPIO General purpose I/O.

H

host Host CPU board. In the Netra CT 810 server, the host CPU board resides in slot 1. In the Netra CT 410, the host CPU board resides in slot 3.

I

IM Information model

IPMI Intelligent Platform Management Interface, used as a communication channel over the cPCI backplane in the Netra CT server.

L

latest occurrence log A managed entity used to group multiple log records to form a latest occurrence log. If no other log record contained in the Latest Occurrence Log instance has values of the attributes identified by the Key Attribute List attribute equal to the attribute values of the log record to be added, the log record is created and contained in the Latest Occurrence Log.

M

MCNet A communication channel running over the CompactPCI backplane. It can be used to communicate between the alarm card, the host CPU board, and any satellite boards.

- MIB** Managed information base used to describe the exchange of information across the network element (NE) interface. A MIB is loadable, but can reference other MIBs.
- module** Software modules are part of a program that are not combined with other parts until the program is linked. Modules do not have to be changed when a new type of object is added.
- MOH** Managed Object Hierarchy. An application that monitors the field replaceable units in the system. MOH runs on the alarm card, the host CPU, and any satellite CPUs.
-

N

- NE** Network Element Managed Entity. A component of the MIB. An instance of this managed entity is automatically created upon initialization
- NFS** Network File System.
- NIS** Network Information System.
-

P

- physical path termination point** See Termination Point MBean.
- PICL** Platform Information and Control Library. A Solaris OS library that provides a method used to publish platform-specific information for clients to access in a way that is not specific to the platform.
- plug-in unit** A managed entity used to represent equipment that is inserted (plugged into) and removed from slots of the NE.
- PMS** Processor Management Service. Manages processor elements used by client applications to implement high availability.

R

- RCM** Reconfiguration Coordination Manager. Part of the Solaris OS's dynamic reconfiguration (DR) framework that enables automated DR removal operations on platforms with appropriate software and hardware configuration.
- RDHCP** Reliable Dynamic Host Configuration Protocol.
- RMI** Remote Method Invocation. Java RMI is a mechanism that allows one to invoke a method on an object that exists in another address space.
- RNFS** Reliable Network File System.

S

- SAT** Satellite. An auxiliary CPU board that occupies a designated cPCI slot on the Netra CT system, which might, under certain conditions, operate independently.
- SNMP** Simple Network Management Protocol. A protocol that allows devices to be controlled remotely by a network management station.
- SMI** Structure of Management Information. A definition that describes the syntax and basic data types available in a given MIB.
- software MBean** A managed entity representing logical information stored in equipment, including programs and data tables. Instances of this managed entity are created by the NE to report to the management system, the currently installed software in the related entity (that is, NE, equipment or Plug-In Unit).
- state change record** A managed entity used to represent logged information resulting from state change notifications. Instances of this managed entity are created automatically by the NE, and deleted by the NE or by request of the managing system.

T

TFTP Trivial File Transfer Protocol.

termination point

MBean A managed entity used to represent the points in the NE where physical paths terminate (such as ports), and physical path level functions (for example, path overhead functions) are performed.

**topology change
notification**

An abstract class representing generic notifications for a change in the topology of a network entity.

Index

A

- access rights, 51
- addressable objects, 51
- agent
 - connecting client, 26
 - netract, 22
- agent, element management, 10, 22, 41
- alarm card, 2, 12, 34
- alarm card view of system, 12, 13, 18
- alarm card, indicating mode, 154
- Alarm Forwarding Discriminator, 62
- alarm pins, 34
- Alarm Severity Trap, 63
- AlarmNotification, 30
- AlarmNotification, example, 30
- alarms
 - assign to objects, example, 35
 - clearing, 36
 - high temperature example, 31
 - set with SNMP, 69
 - setting, 35
- alarms, managing, 30
- AlarmSeverityProfile, example, 31
- assign alarm profile to object, 35
- audience, xiii

B

- Backed Up Alarm Trap, 63
- beginning an application, 21
- board resource management, 80

C

- card, alarm, 2, 12, 34
- cfgadm, 150
- change locationName, 66
- ChassisType, PICL, 141
- Command Line Interface, example, 154
- community strings, 51
- ConditionTime, PICL, 145
- configuration administration, 150
- connecting an agent with a client, example, 26
- ContainmentTreeMBean, example, 27
- CPU cards, managing, 81

D

- DaemonList example, 39
- determining system configuration hierarchy,
 - example, 25
- documentation, related, xv
- DR *see* dynamic reconfiguration
- drawer, definition of, 82
- drivers
 - MCNet, 7
- driving alarm output, 34
- dynamic reconfiguration, 148

E

- element management agent, 22, 41
- ENTITY-MIB, 52
- entPhysicalClass, 53
- entPhysicalContainedIn, 53

- entPhysicalDescr, example, 66
- entPhysicalIndex, 53
- entPhysicalTable, 53
- environment, 1
- example
 - AlarmNotification, 30
 - AlarmSeverityProfile, 31
 - connecting client with agent, 26
 - daemonList, 38
 - finding the root MBean, 27
 - getting nodes on tree, 28
 - initializing PMS client, 85
 - message handling, PMS client, 93
 - monitoring software events, 37
 - NotificationListener, 29
 - PMS client node interface, 124
 - PMS client RND interface, 131
 - PMS client scheduling, 106
 - setting alarm severity with SNMP, 70
 - setting alarms, 34
 - setting watchdog timer, 153
 - SNMP midplane object index, 66
 - system configuration hierarchy, 25
- example, Netra CT security, 23

F

- finding the root MBean, example, 27
- front-access diskless system view, 18
- fru class, PICL, 141
- fru state, PICL, 141
- FRU-ID, changing, 66
- frutree topology, PICL, 140

G

- GeoAddr, PICL, 145
- getting started, 21

H

- hardware
 - associating alarms to failure, 36
- hardware description, 2
- high temperature alarm, SNMP, 69
- HIGH_MEMORY_UTILIZATION, example, 35
- HIGH_TEMPERATURE, example, 35
- host CPU board description, 2
- host CPU board view of system, 13 to 14

- hot-swap, 3

I

- I/O board, description, 3
- initializing PMS client, 85
- instance specifier, 52
- interface
 - dynamic reconfiguration, 148
 - MCNet, 5
 - PMS client node, example, 124
 - PMS client RND, example, 131

J

- Java Dynamic Management Kit
 - see* JDMK
- JDMK
 - agent, 42
 - resources, 44

L

- libcfgadm, 148

M

- managed device, 51
- Managed Object Hierarchy *see* MOH, 6
- managed objects, 10
 - list, 7
- management agent, 22
- Management Information Base *see also* MIB, 50
- managing CPU boards, 81
- MBean
 - introduction to, 42
- MCNet, 159
 - definition, 5
 - description, 7
- memory use alarm tutorial, 34
- message handling, PMS client, example, 93
- MIB
 - access rights, 51
 - addressable objects, 51
 - objects, 50
 - table definition, 52
 - tables, 51
- MIB Notifications, 63
- MIB tables, 50
- midplane FRU-ID, changing, 66

- midplane object
 - sample, 66
- MOH
 - directory path, 44
 - example with SNMP, 66
 - introduction to agent, 22
 - overview, 6
- N**
 - netract, 7
 - netract agent, 22
 - netraCtAlarmSevProfileTable, entry example, 69
 - netraCtHighTempAlarm, example, 70
 - network protocol, 50
 - nodes, example of finding, 28
 - Notification
 - MIB, 63
 - registering a listener, example, 29
 - NotificationFilter, example, 29
 - NotificationListener, example, 29
 - NotificationListener, example, 30
- O**
 - OID (Object Identifiers), 51
 - output alarms, 34
- P**
 - Physical Entity Table, 53
 - physical properties in MIB, 50
 - PICL
 - ChassisType property, 141
 - ConditionTime, 145
 - fru class property, 141
 - Frutree topology, 140
 - GeoAddr, 145
 - man pages, 146
 - port node properties, 142
 - StatusTime, 145
 - temperature sensor node, 146
 - watchdog plug-in, 153
 - PMS, 80 to 135
 - PMS client
 - asynchronous message handling, example, 93
 - initializing, example, 85
 - RND interface, example, 131
 - scheduling, example, 106
 - PMS client node interface, example, 124
 - PMS introduction, 6
 - PMS software, overview, 80
 - port class, PICL, 142
 - port condition, PICL, 144
 - port state, PICL, 143
 - portType, PICL, 144
 - processor management services, 80 to 135
 - processor management services *see also* PMS, 6
- R**
 - RCM, 148
 - rear-access diskfull system view, 17, 19
 - rear-access diskless system view, 18
 - Reconfiguration Coordination Manager, 148
 - registering notification listener, example, 29
 - Remote Method Invocation (RMI), 8, 43
 - represent the system MBeans, example, 25
 - RFC2578, 51
 - RFC2579, 51
 - RFC2737, 52
 - RG (Resource Group) description, 81
 - RMI API directory path, 44
 - RMI *see* Remote Method Invocation
 - root MBean, example of finding, 27
 - routing tables, in MIB, 50
- S**
 - satellite CPU board, 3, 15
 - satellite CPU board rear-access view, 20
 - set alarms with SNMP, 69
 - setting watchdog timer, 153
 - SNMP
 - setting high temperature alarm, example, 70
 - SNMP interface, 7
 - SNMP Traps, 50, 62 to 63
 - software environment, 1
 - software service daemons, example, 38
 - starting netract agent, 21
 - StatusTime, PICL, 145
 - system configuration hierarchy example, 25
 - system view
 - from alarm card, 12, ?? to 13
 - from host CPU, 13 to 14

- from satellite CPU, 15
- front-access, ?? to 13
- front-access diskfull, 18
- rear-access, 13, 14, 23
- rear-access diskfull, 17, 19
- satellite CPU board, 20

T

- table definition, 52
- tables in MIB, 50
- temperature
 - sensor node, PICL, 146
- temperature alarm tutorial, 34
- thermister, 35
- timer, watchdog, 153
- Trap
 - Alarm Backed Up, 63
 - Alarm Severity, 63
 - definition, 50
- tutorial, 21

W

- watchdog plug-ins, 155
- watchdog timer, 153
- watchdog-timer settings, 156