



System Management Services (SMS) 1.3 Dynamic Reconfiguration User Guide

Creating System Management Applications on the
Sun Fire™ 15K and Sun Fire 12K Systems

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.

Part No. 816-7723-10
January 2003 Revision A

Send comments about this document to: docfeedback@sun.com

Copyright 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or 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 other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Fire, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Sun Fire, OpenBoot, Sun Management Center, Sun RSM Array, JavaHelp, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and 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.

Federal Acquisitions: Commercial Software—Government Users Subject to Standard License Terms and Conditions.

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 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 Etats-Unis. Tous droits réservés.

Ce produit ou document est 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 Fire, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Sun Fire, OpenBoot, Sun Management Center, Sun RSM Array, JavaHelp, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, 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 licenciés 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'ETAT" ET TOUTES AUTRES CONDITIONS, DECLARATIONS ET GARANTIES EXPRESSES OU TACITES SONT FORMELLEMENT EXCLUES, DANS LA MESURE AUTORISEE PAR LA LOI APPLICABLE, Y COMPRIS NOTAMMENT TOUTE GARANTIE IMPLICITE RELATIVE A LA QUALITE MARCHANDE, A L'APTITUDE A UNE UTILISATION PARTICULIERE OU A L'ABSENCE DE CONTREFAÇON.



Adobe PostScript

Contents

Preface vii

Before You Read This Book vii

How This Book Is Organized vii

Using UNIX Commands viii

Typographic Conventions ix

Shell Prompts ix

Related Documentation x

Accessing Sun Documentation Online x

Sun Welcomes Your Comments x

1. Introduction to DR in System Management Services 1

What Is DR? 1

 Automatic DR 1

 Enhanced System Availability 2

Component Types 2

DR on I/O Boards 3

 DR on hsPCI+ I/O Boards 3

 Golden IOSRAM 3

Capacity on Demand (COD) 4

 DR on COD Boards 4

Sun Fire 15K/12K Domains	4
Enabling DR on Domains Running the Solaris 8 2/02 Operating Environment	5
DR Administration Models	6
DR Software Components on the SC	6
Domain Configuration Agent (DCA)	6
Platform Configuration Daemon (PCD)	7
Domain X Server (DXS)	7

2. SMS DR Procedures 9

Showing Device Information	9
▼ To Show Device Information	10
Showing Platform Information	11
▼ To Show Platform Information	12
Showing Board Information	12
SC State Models	12
The showboards(1M) command	13
▼ To Show Board Information	14
Adding Boards	14
▼ To Add a Board to a Domain	15
Deleting Boards	15
▼ To Delete a Board From a Domain	16
Moving Boards	16
▼ To Move a Board	17
Replacing Active System Boards	17
▼ To Replace an Active System Board	17

3. SMS DR User Interfaces 19

SMS DR Commands and Options	19
addboard(1M)	19

deleteboard(1M)	21
moveboard(1M)	22
rcfgadm(1M)	25
scdrhelpl(1M)	26
showboards(1M)	26
showdevices(1M)	27
showplatform(1M)	28
Error Message Help System	28
JavaHelp Table of Contents	29
JavaHelp Index	29
JavaHelp Search	30

Preface

This book describes how you use the dynamic reconfiguration (DR) feature by entering commands in the system management services (SMS) software. DR enables you to attach system boards to, and detach them from, active Sun Fire™ 15K and Sun Fire 12K domains while the Solaris™ operating environment continues to run.

Before You Read This Book

This book is intended for the Sun Fire 15K/12K platform administrator who has a working knowledge of UNIX® systems, particularly those based on the Solaris operating environment. If you do not have such knowledge, first read the Solaris user and system administrator books provided with this system, and consider UNIX system administration training.

How This Book Is Organized

This book contains the following chapters:

Chapter 1 “Introduction to DR in System Management Services”

Chapter 2 “SMS DR Procedures”

Chapter 3 “SMS DR User Interfaces”

Using UNIX Commands

This document does not contain complete information on basic UNIX® commands and procedures, such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- Online documentation for the Solaris™ software environment
- Other software documentation that you received with your system

Typographic Conventions

Typeface or Symbol	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	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.
	Command-line variable; replace with a real name or value	To delete a file, type <code>rm filename</code> .

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

Application	Title	Part Number
DR User Guide	<i>Sun Fire 15K/12K Dynamic Reconfiguration User Guide</i>	816-5075
Platform-specific release notes	<i>Solaris 9 4/03 Release Notes Supplement for Sun Hardware</i>	817-1106
SMS Administrator Guide	<i>System Management Services (SMS) 1.3 Administrator Guide</i>	816-5319
SMS Release Notes	<i>System Management Services (SMS) 1.3 Release Notes</i>	816-5321
DR Webpage	http://www.sun.com/servers/highend/dr_sunfire	n/a

Accessing Sun Documentation Online

You can view and print a broad selection of Sun(TM) documentation, including localized versions, at:

<http://www.sun.com/documentation>

You can also purchase printed copies of select Sun documentation from iUniverse, the Sun documentation provider, at:

<http://corppub.iuniverse.com/marketplace/sun>

Sun Welcomes Your Comments

Sun is interested in improving its documentation and welcomes your comments and suggestions. You can email your comments to Sun at:

docfeedback@sun.com

Please include the part number of this document (816-7723-10) in the subject line of your email.

Introduction to DR in System Management Services

This chapter contains an introduction to the dynamic reconfiguration (DR) feature in the system management services (SMS) software on the Sun Fire 15K/12K system controller (SC).

What Is DR?

The dynamic reconfiguration feature on the Sun Fire 15K/12K server enables you to perform hardware configuration changes to a live domain that is running the Solaris operating environment, without causing machine downtime. You can also use DR, in conjunction with hot-swap functionality, to physically remove boards from, or add them to, the server.

You can execute DR operations from the SC by using the system management services commands—`addboard(1M)`, `moveboard(1M)`, `deleteboard(1M)`, and `rcfgadm(1M)`.

Note – You can execute DR operations either on the SC, or on the domain, using the `cfgadm(1M)` command. Refer to the *Sun Fire 15K/12K Dynamic Reconfiguration User Guide* for more information about running DR on the domain.

Automatic DR

Automatic DR enables an application to execute DR operations without requiring user interaction. This ability is provided by an enhanced DR framework that includes the reconfiguration coordination manager (RCM) and the system event

facility, `sysevent`. The RCM enables application-specific loadable modules to register callbacks. The callbacks perform preparatory tasks before a DR operation; error recovery during a DR operation; or clean-up after a DR operation. The system event framework enables applications to register for system events and receive notifications of those events. The automatic DR framework interfaces with the RCM and with the system event facility to enable applications to automatically give up resources prior to unconfiguring them, and to capture new resources as they are configured into the domain.

The automatic DR framework can be used both locally (that is, from the domain by using the `cfgadm(1M)` command) or from the SC. The automatic DR operations that are initiated locally on the domain are referred to as *local* automatic DR, and the automatic DR operations initiated from the SC are referred to as *global* automatic DR. The global automatic DR operations include moving system boards from one domain to another, configuring hot-swapped boards into a domain, and removing system boards from a domain.

Enhanced System Availability

The DR feature enables you to hot-swap system boards without bringing the server down. It is used to unconfigure the resources on a faulty system board from a domain so that the system board can be removed from the server. The repaired, or replacement, board can be inserted into the domain while the Solaris operating environment is running. DR then configures the resources on the board into the domain. If you use the DR feature to add or remove a system board or component, DR always leaves the board or component in a known configuration state (see the section “SC State Models” on page 12 for more information about configuration states for system boards and components).

Component Types

You can use DR to add or to remove the following components:

Component	Description
<code>cpu</code>	An individual CPU
<code>memory</code>	All of the memory on the board
<code>pci</code>	Any I/O device, controller, or bus

DR on I/O Boards

You must use caution when you add or remove system boards with I/O devices. Before you can remove a board with I/O devices, all of its devices must be closed and all its file systems must be unmounted.

If you need to remove a board with I/O devices from a domain temporarily and then re-add it before any other boards with I/O devices are added, reconfiguration is not necessary and need not be performed. In this case, device paths to the board devices will remain unchanged. But if you add another board with I/O devices after the first was removed and then re-add the first board, reconfiguration is required because the paths to devices on the first board have changed.

Note – Before attempting to perform DR operations on an I/O board in a domain, make sure at least two CPUs are available to the domain. Further, make sure at least one of those CPUs is located on a CPU/memory board, and that no processes are bound to it. See the `pbind(1M)` man page for more information about bound processes.

DR on hsPCI+ I/O Boards

DR supports dynamic reconfiguration of hsPCI+ I/O boards. Each hsPCI+ I/O board includes two XMITs ASICs and four hot-pluggable hsPCI slots.

Golden IOSRAM

Each I/O board in a domain contains an IOSRAM device. However, only one IOSRAM device, called the *golden IOSRAM*, is used for SC-to-domain communications at a time. The golden IOSRAM contains the “tunnel” that is used for SC-to-domain communications. Because DR can remove I/O boards, it is sometimes necessary to stop using the current golden IOSRAM and make another IOSRAM device the golden IOSRAM. This process is called a “tunnel switch,” and takes place whenever DR unconfigures the current golden IOSRAM.

When a domain is booted, the lowest-numbered I/O board in the domain is typically selected to be the initial golden IOSRAM.

Capacity on Demand (COD)

The COD option provides additional CPU resources on COD CPU/Memory boards that you install in your Sun Fire 15K/12K system. Although your Sun Fire 15K/12K system comes configured with a minimum number of standard (active) CPU/Memory boards, it can have a mix of both standard and COD CPU/Memory boards installed, up to a maximum 18 boards. At least one active CPU is required for each domain in the system.

DR on COD Boards

You can use DR to move COD boards into and out of domains in the same way you use DR to move standard CPU/memory boards.

You can use the CPUs on a COD board only after you purchase right-to-use (RTU) licenses for them. Each COD RTU license entitles you to receive a COD RTU license key that enables a specified number of CPUs on COD boards in a single system. Whenever you use DR to configure a COD board into a domain, make sure that enough RTU licenses are available to the target domain to enable each active CPU on the COD board. If there are not enough RTU licenses available to a target domain when you add a COD board, a status message is displayed for each CPU that cannot be enabled in the domain.

For more information about the COD option, see the System Management Services (SMS) 1.3 *Administrator Guide*.

Sun Fire 15K/12K Domains

The Sun Fire 15K/12K server can be divided into domains. Each domain is based on the system board slots that are assigned to it. Further, each domain is electrically isolated into hardware partitions, which ensures that any failure in one domain does not affect the other domains in the server.

Sun Fire 15K/12K domain configuration is determined by the domain configuration in the platform configuration database (PCD), which resides on the SC. The PCD controls how the system board slots are logically partitioned into domains. The domain configuration represents the intended domain configuration. Thus, the configuration can include empty slots and populated slots. The physical domain is determined by the logical domain.

The number of slots available to a given domain is controlled by an available component list (ACL) that is maintained on the SC. A slot must be assigned or available to a domain before you can change its state. After a slot has been assigned to a domain, it becomes visible to that domain and unavailable and invisible to any other domain. Conversely, you must disconnect and unassign a slot from its domain before you can assign and connect it to another domain.

The logical domain is the set of slots that belong to the domain. The physical domain is the set of boards that are physically interconnected. A slot can be a member of a logical domain without having to be part of a physical domain. After the domain is booted, the system boards and the empty slots can be assigned to or unassigned from a logical domain; however, they are not allowed to become a part of the physical domain until the operating system requests it. System boards or slots that are not assigned to any domain are available to all domains. These boards can be assigned to a domain by the platform administrator; however, an available component list can be set up on the SC to allow users with appropriate privileges to assign available boards to a domain.

Enabling DR on Domains Running the Solaris 8 2/02 Operating Environment

While the Solaris 9 operating environment supports the full functionality of DR, Solaris 8 2/02 was the first release of Solaris operating environment to support DR of I/O boards.

You can enable the full functionality of DR on domains running a version of the operating environment no earlier than Solaris 8 2/02 by installing patches and a new kernel update on the domain; and by installing the System Management Services (SMS) 1.3 software on the system controller (SC).

For complete information and instructions for enabling DR on such a domain, visit:

http://www.sun.com/servers/highend/dr_sunfire

DR Administration Models

The available component list controls what administrative tasks can be performed, based on the name and group identification of the user. A brief description of the privileges model for each DR operation is given in Chapter 3 “*SMS DR User Interfaces*”. For a detailed description of the privileges required for each SMS command, refer to the *System Management Services (SMS) 1.3 Administrator Guide*.

DR Software Components on the SC

Various processes and daemons on the Sun Fire 15K/12K system controller (SC) work together to accomplish DR operations. The processes and/or daemons that are used depends entirely on the point of execution of the DR operation. For instance, if you execute the DR operation from the SC, the system uses several more processes and/or daemons to accomplish the DR operation than it would if you executed the DR operation from the domain.

For more information about the processes and daemons that reside on the domain, refer to the *Sun Fire 15K/12K Dynamic Reconfiguration User Guide*. In addition, refer to the *System Management Services (SMS) 1.3 Administrator Guide* for more information about the processes and daemons that reside in the SMS software on the SC.

Domain Configuration Agent (DCA)

The domain configuration agent (DCA) enables applications such as Sun Management Center and SMS to initiate DR operations on a Sun Fire 15K/12K domain. The DCA runs on the SC and manages the DR communications between software applications running on the SC and the domain configuration server on the domain. An individual instance of the DCA runs on the SC for each domain on the Sun Fire 15K/12K system. For more information about the DCA, refer to the *System Management Services (SMS) 1.3 Administrator Guide*.

Note – If you alter or remove the `sun-dr` entry in the `inetd.conf` file, make the same change to the `sun-dr` entry in the `ipsecinit.conf` file.

Platform Configuration Daemon (PCD)

The platform configuration daemon (PCD) manages the configuration of the Sun Fire 15K/12K systems through a collection of flat files that comprise the PCD database. All changes to the configuration of the Sun Fire 15K/12K system must go through the PCD. For more information about the PCD, refer to the *System Management Services (SMS) 1.3 Administrator Guide*.

Domain X Server (DXS)

The domain x server (DXS) manages communication between the SC and the DR module (`drmach`) on the domain. An individual instance of the DXS runs on the SC for each domain on the Sun Fire 15K/12K system. For more information about the DCX, refer to the *System Management Services (SMS) 1.3 Administrator Guide*.

SMS DR Procedures

This chapter contains procedures that describe how to use the DR feature on the Sun Fire 15K/12K system controller (SC), which runs the system management services (SMS) software. The following procedures are included:

- “Showing Device Information” on page 9
- “Showing Board Information” on page 12
- “Showing Platform Information” on page 11
- “Adding Boards” on page 14
- “Deleting Boards” on page 15
- “Moving Boards” on page 16
- “Replacing Active System Boards” on page 17

Showing Device Information

Before you attempt to perform any DR operation, use the `showdevices(1M)` command to display the device information, especially when you are removing devices.

▼ To Show Device Information

- Use the `showdevices(1M)` command to display the device information for the domain.

```
% showdevices -v -d A
```

The above command displays the device information for all of the devices in the domain and produces output such as the following:

```
CPU
----
domain  board  id   state   speed  ecache  usage
A       SB1    40   online  400    4
A       SB1    41   online  400    4
A       SB1    42   online  400    4
A       SB1    43   online  400    4
A       SB2    55   online  400    4
A       SB2    56   online  400    4
A       SB2    57   online  400    4
A       SB2    58   online  400    4
```

The following output represents an example of the memory output for the `showdevices(1M)` command above.

```
Memory
drain in progress:
-----
          board  perm   base   domain  target  deleted  remaining
domain  board  mem MB  mem MB  addr   mem MB  board  MB      MB
A       SB1    2048   933    0x600000 4096   C2     250    1500
A       SB2    2048   0      0x200000 4096
```

The following output represents an example of the I/O devices output for the `showdevices(1M)` command above.

```
IO Devices
-----
domain  board  device  resource          usage
A       I01    sd0
A       I01    sd1
A       I01    sd2
A       I01    sd3      /dev/dsk/c0t3d0s0  mounted filesystem "/"
A       I01    sd3      /dev/dsk/c0t3d0s1  dump device (swap)
A       I01    sd3      /dev/dsk/c0t3d0s1  swap area
A       I01    sd3      /dev/dsk/c0t3d0s3  mounted filesystem "/var"
A       I01    sd3      /var/run           mounted filesystem "/var/run"
A       I01    sd4
A       I01    sd5
```

Refer to the section “`showdevices`” on page 27 for more information about the `showdevices(1M)` command. Refer to the `showdevices(1M)` man page for a complete list of the options and arguments for this command, and for more information about displaying device-specific information.

Showing Platform Information

Before you attempt to add, move, or delete a board to or from a specific domain, use the `showboards(1M)` command to determine the domain ID, the boards available to the domain, and the status of the domain.

You can use the domain ID in all of the DR commands. You can use the board list to determine the domain to which a specific board is assigned, and you can use the domain status to determine whether or not you can add, delete, or move a board to or from the domain. Use the `showplatforms(1M)` command to determine whether the component is in the available component list.

You must have the appropriate privileges to use the `showplatforms(1M)` command. See the section “`showplatform`” on page 28 for more information about the `showplatforms(1M)` command, including a table showing which user groups can use it. Also, refer to the `showplatform(1M)` man page for more information.

▼ To Show Platform Information

- Use the `showplatform(1M)` command to obtain the domain information.

```
% showplatform
```

The `showplatform(1M)` command displays the domain ID, the available component list, and the status of the domain, as in the following example.

```
ACLs for domain domainA:
  slot0: SB0, SB1, SB2, SB3
  slot1: IO0, IO1, IO2, IO3

ACLs for domain domainB:
  slot0: None
  slot1: None

Domain          Solaris Nodename      Domain Status
domainA         sms3-b0               Powered Off
domainB         sms3-b1               Running Solaris
```

Showing Board Information

Before you attempt to delete or move a system board, you must query the board to determine the state of the board and the domain to which it is assigned.

SC State Models

On the Sun Fire 15K/12K SC, a board can be in one of four states: unavailable, available, assigned, or active.

Note – The state of a board on the SC is not the same as the state of a board on the domain. For more information about board states on the domain, refer to the *Sun Fire 15K/12K Dynamic Reconfiguration User Guide*.

The following table contains descriptions of board states on the SC.

Name	Description
unavailable	The board is unavailable to the domain. The board has not been added to the available component list for the specified domain, or the board is currently assigned to another domain. Note that boards that are not in the available component list are invisible to the domain. In the <code>unavailable</code> state, the board is not considered part of the specified domain.
available	The board is available to be added to the domain. The board is in the available component list for the domain. Note that the board can be available to any number of domains. In the <code>available</code> state, the board is not considered to be part of the logical domain.
assigned	The board has been assigned to the domain, and may be in the domain's available component list. The board is unavailable to any other domain. In the <code>assigned</code> state, the board is considered to be part of the logical domain.
active	The board has been connected. Or, the board has been connected and configured into the Solaris operating environment and is available for use by the operating system. In the <code>active</code> state, the board is considered part of the physical domain.

The `showboards(1M)` command

After you have determined the domain ID that contains the board that you want to delete or move, or after you have determined that a particular board has already been assigned to a specific domain, use the `showboards(1M)` command to determine the state of the board. The board may be in a state that makes it impossible for you to delete or move it.

Note – The output of the `showboards(1M)` command depends on the privileges of the user. For instance, the platform administrator can obtain information about all of the boards in the server. The domain administrator and domain configurator, however, can obtain the information about only those boards that are assigned and available to the domain(s) to which they have access. Refer to the section “`showboards`” on page 26 for more information.

▼ To Show Board Information

- Use the `showboards(1M)` command to display the board information for the domain.

```
% showboards -d A
```

The above command displays the device information for domain A. The following example contains the information that would be displayed.

Slot	Pwr	Type of Board	Board Status	Test Status	Domain
SB0	On	CPU Board	Active	Passed	A
SB1	-	Empty Slot	Assigned	-	A

You can use the `showboards(1M)` command to display all of the assigned system boards, all of the available system boards, and all of the CPU/memory and I/O boards in the domain. Refer to the `showboards(1M)` man page for more information about showing board information.

Adding Boards

Adding a board to a domain moves the board through several state changes. If it is not already assigned, it is first assigned to the domain. Then, it is connected to the domain and configured into the Solaris operating environment. After it is connected, it is considered part of the physical domain and available for use by the operating system.

You must have the appropriate privileges to add a board to a domain. See the section “`addboard`” on page 19 for more information, including a description of the privileges needed to use this command.

Note – Whenever you use DR to add a COD board into a domain, make sure that enough RTU licenses are available to the target domain to enable each active CPU on the COD board. If there are not enough RTU licenses available to a target domain when you add a COD board to it, DR displays a message for each CPU that cannot be enabled in the domain. For more information about the COD option, see the System Management Services (SMS) 1.3 *Administrator Guide*.

▼ To Add a Board to a Domain

- Use the `addboard(1M)` command to add the board to the domain.

The following example of the `addboard(1M)` command adds system board 2 (SB2) to the domain specified by *domain_id*. Two retries are performed, if necessary, with a wait time of 10 minutes (600 seconds) between retries.

```
% addboard -d domain_id -r 2 -t 600 SB2
```

Note – If the `addboard(1M)` command fails during a DR operation, the board does not return to its original state. A `dxs` or `dca` error message is logged to the domain. If the error is recoverable, you can retry the command. If the error is unrecoverable, you will need to reboot the domain to use the board.

Deleting Boards

Deleting a board from a domain removes the board from the domain to which it is currently assigned, and in which it may be active. To delete a board, it must be in the assigned or active state.

Always check the usage of the components on a board before you delete it from a domain. If the board hosts permanent memory, the memory is moved to another board within the same domain before the board is deleted from the domain. Likewise, if any busy devices are present, you must wait or ensure that the device is no longer being used by the system before you attempt to remove the board.

A domain administrator can unconfigure and disconnect a board, but cannot unassign a board from the a domain unless the board is in the available component list. See the section “`deleteboard`” on page 21 for more information, including a description of privileges required to use this command.

▼ To Delete a Board From a Domain

- Use the `deleteboard(1M)` command to delete the board from the domain.

The following example of the `deleteboard(1M)` command deletes system board 2 (SB2) from its current domain. Two retries are performed, if necessary, with a wait time of 15 minutes (900 seconds) between retries.

```
% deleteboard -r 2 -t 900 SB2
```

Note – If the `deleteboard(1M)` command fails during a DR operation, the board does not return to its original state. A `dxs` or `dca` error message is logged to the domain. If the error is recoverable, you can retry the command. If the error is unrecoverable, you will need to reboot the domain to use the board.

Moving Boards

Moving a board from one domain to another domain is performed in several steps. First, the board is removed from the domain to which it is currently assigned, and in which it may be active; the board must be in the assigned or active state. Next, it is assigned to the target domain. Then, it is connected to the target domain and configured into the Solaris operating environment, where it becomes available for use.

You should always check the usage of the memory and devices on a board before you move it out of a domain. If the board hosts permanent memory, the memory must be moved to another board within the same domain before the board can be moved to another domain. Likewise, if any busy devices are present, you must wait or ensure that the device is no longer being used by the system before you attempt to move the board.

See the section “`moveboard`” on page 22 for more information, including a description of privileges required to use the `moveboard(1M)` command.

Note – Whenever you use DR to move a COD board into a domain, make sure that enough RTU licenses are available to the target domain to enable each active CPU on the COD board. If there are not enough RTU licenses available to a target domain when you add a COD board to it, DR displays a message for each CPU that cannot be enabled in the domain. For more information about the COD option, see the System Management Services (SMS) 1.3 *Administrator Guide*.

▼ To Move a Board

- Use the `moveboard(1M)` command to move the board from one domain to another domain.

The following example of the `moveboard(1M)` command moves system board 2 (SB2) from its current domain to the domain specified by *domain_id*. Two retries are performed, if necessary, with a wait time of 15 minutes (900 seconds) between retries.

```
% moveboard -d domain_id -r 2 -t 900 SB2
```

Note – If the `moveboard(1M)` command fails during a DR operation, the board does not return to its original state. A `dxs` or `dca` error message is logged to the domain. If the error is recoverable, you can retry the command. If the error is unrecoverable, you will need to reboot the domain to use the board.

Replacing Active System Boards

This section describes how to replace a system board that is active in a domain with another system board.

▼ To Replace an Active System Board

In the following steps, system board 2 (SB2) is removed from its current domain and replaced by system board 3 (SB3).

1. Use the `deleteboard(1M)` command to delete system board 2 (SB2) from the domain.

This step deletes system board 2 from its current domain. Two retries are performed, if necessary, with a wait time of 15 minutes (900 seconds) between retries.

```
% deleteboard -r 2 -t 900 SB2
```

2. Use the `addboard(1M)` command to add system board 3 (SB3) to the domain.

This step adds system board 3 to the domain specified by *domain_id*. Two retries are performed, if necessary, with a wait time of 15 minutes (900 seconds) between retries.

```
% addboard -d domain_id -r 2 -t 900 SB3
```

SMS DR User Interfaces

The dynamic reconfiguration (DR) feature supports a single point of control, the Sun Fire 15K/12K system controller (SC), which runs the system management services (SMS) software. However, DR also supports administration control from the domain. This chapter describes the DR interfaces on the SC. For information about using `cfgadm(1M)` to run DR on the domain, refer to the *Sun Fire 15K/12K Dynamic Reconfiguration User Guide*.

SMS DR Commands and Options

This section contains descriptions of the SMS DR commands and related options. For more information about each SMS DR command, refer to the *System Management Services (SMS) 1.3 Reference Manual*.

`addboard(1M)`

The `addboard(1M)` command attaches a specific board to a specific domain. Refer to the section “Adding Boards” on page 14, and to the `addboard(1M)` man page, for more information.

The following table describes the `addboard(1M)` command options and operands:

Options and Operands	Specifies
<code>board_id</code>	The board ID of the board to be added. The board ID corresponds to the board location (for example, SB2 is the board in slot 2). Multiple board identifiers are permitted.
<code>-c function</code>	The configuration state into which to move the board. You can add a board by steps. For example, you can assign the board, connect it, then configure it.
<code>-d domain_id domain_tag</code>	Identifies the target domain.
<code>-f</code>	Forces the specified action to occur. Typically, this is a hardware-specific override of a safety feature. Forcing a state change operation can allow use of the hardware resources of an occupant that is not in the <code>ok</code> or <code>unknown</code> conditions, at the discretion of any hardware-dependent safety checks.
<code>-h</code>	Help, which displays the usage description
<code>-n</code>	No to all prompts
<code>-q</code>	Quiet mode, which means messages are not written to standard output
<code>-r retry_count</code>	The number of retries that are performed if the operation fails
<code>-t timeout</code>	Wait time in seconds between retries
<code>-y</code>	Yes to all prompts

The following table contains the privileges needed to use the `addboard(1M)` command. The platform operator, platform service, and superuser groups cannot initiate the `addboard(1M)` command.

Platform Admin	Domain Admin	Domain Configurator
Can assign boards to a domain by using the <code>-c</code> option with the <code>assign</code> function.	Can connect or configure a board into the domain if the board has been assigned to the domain, or if it appears in the available component list for the domain and is not assigned to another domain.	Can connect or configure a board into the domain if the board has been assigned to the domain, or if it appears in the available component list for the domain and is not assigned to another domain.

The following example of the `addboard(1M)` command attaches system board 2 (SB2) to the domain specified by *domain_id*. Two retries are performed, if necessary, with a wait time of 10 minutes (600 seconds) between retries.

```
% addboard -d domain_id -r 2 -t 600 SB2
```

Note – If the `addboard(1M)` command fails during a DR operation, the board does not return to its original state. A `dxs` or `dca` error message is logged to the domain. If the error is recoverable, you can retry the command. If the error is unrecoverable, you will need to reboot the domain to use the board.

deleteboard(1M)

The `deleteboard(1M)` command attempts to detach a specified board from the domain in which it currently resides. Refer to the section “Deleting Boards” on page 15, and to the `deleteboard(1M)` man page, for more information.

The following table describes the `deleteboard(1M)` command options and operands:

Options and Operands	Specifies
<i>board_id</i>	The board ID of the board to be deleted. The board ID corresponds to the board location (for example, SB2 is the system board in slot 2). Multiple board identifiers are permitted.
-c <i>function</i>	The configuration state into which to move the board. You can delete a board by steps. For example, you can unconfigure the board, disconnect it, and then unassign it.
-f	Forces the specified action to occur. Typically, this is a hardware-specific override of a safety feature. Forcing a state change operation can allow use of the hardware resources of an occupant that is not in the <code>ok</code> or <code>unknown</code> conditions, at the discretion of any hardware-dependent safety checks.
-h	Help, which displays the usage description
-n	No to all prompts
-q	Quiet mode, which means messages are not written to standard output.

Options and Operands	Specifies
<code>-r <i>retry_count</i></code>	Number of retries that are performed if the detach fails
<code>-t <i>timeout</i></code>	Wait time in seconds before a retry
<code>-y</code>	Yes to all prompts

The following table contains the privileges needed to use the `deleteboard(1M)` command. The platform operator, platform service, and superuser groups cannot initiate the `deleteboard(1M)` command.

Platform Admin	Domain Admin	Domain Configurator
Can unassign boards that are not active in a domain by using the <code>-c</code> option with the <code>unassign</code> function. If the user also has domain privileges, <code>deleteboard</code> also unconfigures and disconnects the board before it unassigns it.	Can unconfigure, disconnect or unassign a board from the domain. The board can be unassigned from the domain only if it appears in the available component list.	Can unconfigure, disconnect or unassign a board from the domain. The board can be unassigned from the domain only if it appears in the available component list.

The following example of the `deleteboard(1M)` command detaches system board 2 (SB2) from its current domain. Two retries are performed, if necessary, with a wait time of 15 minutes (900 seconds) between retries.

```
% deleteboard -r 2 -t 900 SB2
```

Note – If the `deleteboard(1M)` command fails during a DR operation, the board does not return to its original state. A `dxs` or `dca` error message is logged to the domain. If the error is recoverable, you can retry the command. If the error is unrecoverable, you will need to reboot the domain to use the board.

moveboard(1M)

The `moveboard(1M)` command detaches a board from the domain in which it currently resides and attaches it to a specified domain. Refer to the section “Moving Boards” on page 16, and to the `moveboard(1M)` man page, for more information.

The following table describes the `moveboard(1M)` command options and operands:

Options and Operands	Specifies
<code>board_id</code>	The board ID of the board to be moved. The board ID corresponds to the board location (for example, SB2 is the system board in slot 2). Multiple board identifiers are permitted.
<code>-c function</code>	The configuration state into which to move the board. You can move a board by steps. For example, you can assign the board, connect it, and then configure it.
<code>-d domain_id domain_tag</code>	Identifies the target domain.
<code>-f</code>	Forces the specified action to occur. Typically, this is a hardware-specific override of a safety feature. Forcing a state change operation can allow use of the hardware resources of an occupant that is not in the <code>ok</code> or <code>unknown</code> conditions, at the discretion of any hardware-dependent safety checks.
<code>-h</code>	Help, which displays the usage description
<code>-n</code>	<i>No</i> to all prompts
<code>-q</code>	Quiet mode, which means messages are not written to standard output
<code>-r retry_count</code>	Number of retries that are performed if the operation fails
<code>-t timeout</code>	Wait time in seconds before a retry
<code>-y</code>	<i>Yes</i> to all prompts

The following table contains the privileges needed to use the `moveboard(1M)` command. The platform operator, platform service, and superuser groups cannot initiate the `moveboard(1M)` command.

Platform Admin	Domain Admin	Domain Configurator
Can re-assign boards from one domain to another domain by using the <code>-c</code> option with the <code>assign</code> function. The board cannot be active in the domain from which it is being re-assigned.	Can assign, connect, or configure a board that is in another domain. If the board is active in another domain, the <code>moveboard</code> command unconfigures and disconnects the board from that domain. The board must be in the available component list in order to unassign and re-assign it using <code>moveboard</code> . The <code>moveboard</code> command can connect and configure the board.	Can assign, connect, or configure a board that is in another domain. If the board is active in another domain, the <code>moveboard</code> command unconfigures and disconnects the board from that domain. The board must be in the available component list in order to unassign and re-assign it using <code>moveboard</code> . The <code>moveboard</code> command can connect and configure the board.
	The domain administrator must have domain privileges for both domains to use the <code>moveboard(1M)</code> command.	The domain configurator must have domain privileges for both domains to use the <code>moveboard(1M)</code> command.

The following example of the `moveboard(1M)` command moves system board 5 (SB5) from its current domain to the domain specified by `domain_id`. Two retries are performed, if necessary, with a wait time of 15 minutes (900 seconds) between retries.

```
% moveboard -d domain_id -r 2 -t 900 SB5
```

Note – If the `moveboard(1M)` command fails during a DR operation, the board does not return to its original state. A `dxs` or `dca` error message is logged to the domain. If the error is recoverable, you can retry the command. If the error is unrecoverable, you will need to reboot the domain to use the board.

rcfgadm(1M)

The `rcfgadm(1M)` command provides remote configuration administration operations on dynamically reconfigurable hardware resources. The `rcfgadm(1M)` command performs configuration administration on attachment points, which are device nodes in the device tree.

The following table describes the `rcfgadm(1M)` command options and operands.

Options and Operands	Specifies
<code>-a</code>	The <code>-a</code> option lists dynamic attachment points.
<code>-c function</code>	The configuration state into which to move the board. You can disconnect, connect, configure, or unconfigure a board.
<code>-d domain_id domain_tag</code>	Identifies the target domain.
<code>-f</code>	The specified action to be forced to occur
<code>-h ap_id ap_type</code>	Help message text to be printed. If <code>ap_id</code> or <code>ap_type</code> is given, the command displays the hardware specific help for the attachment point.
<code>-l ap_id ap_type</code>	State and condition of attachment points to be listed
<code>-n</code>	<i>No</i> to all prompts
<code>-o hardware_options</code>	Hardware specific options
<code>-r retry_count</code>	Number of times to retry the DR operation
<code>-s listing_options</code>	Options to list
<code>-T timeout</code>	Wait time in seconds between retries
<code>-t</code>	A test on one or more attachment points
<code>-v</code>	Specifies verbose mode.
<code>-x hardware_function</code>	Hardware specific functions
<code>-y</code>	<i>Yes</i> to all prompts

Refer to the `rcfgadm(1M)` man page for more information and examples of how to use this command.

The following table contains the privileges needed to use the `rcfgadm(1M)` command. The platform operator, platform service, and superuser groups cannot initiate the `rcfgadm(1M)` command.

Platform Admin	Domain Admin	Domain Configurator
Can assign boards to, or unassign boards them from, a domain by using the <code>-x</code> option with the <code>assign</code> or <code>unassign</code> function, respectively. To use the <code>unassign</code> function, the board must be assigned and cannot be active in a running domain.	Can disconnect, connect, configure, or unconfigure a board to or from the domain. Can assign or unassign a board if the board is in the domain's available component list.	Can disconnect, connect, configure, or unconfigure a board to or from the domain. Can assign or unassign a board if the board is in the domain's available component list.

Note – If the `rcfgadm(1M)` command fails during a DR operation, the board does not return to its original state. A `dxs` or `dca` error message is logged to the domain. If the error is recoverable, you can retry the command. If the error is unrecoverable, you will need to reboot the domain to use the board.

scdrhelp(1M)

The `scdrhelp(1M)` shell script starts the Sun Fire 15K/12K dynamic reconfiguration errors help system. The help system uses the JavaHelp™ `hviewer` script.

All of the user privileges groups can use this command, except for the domain administrator and the domain configurator.

Refer to the section Section “Error Message Help System” on page 328 and to the `scdrhelp(1M)` man page for more information about this script.

showboards(1M)

The `showboards(1M)` command displays the assignment information and status of system boards in a domain; and indicates whether a board is a Capacity On Demand (COD) board. Refer to the section “Showing Board Information” on page 12, and to the `showboards(1M)` man page, for more information.

Although the `showboards(1M)` command is not DR-specific, it should be used in conjunction with the DR commands. The following table describes the `showboards(1M)` command options.

Option	Specifies
<code>-d domain_id domain_tag</code>	Identifies the target domain.
<code>-h</code>	Help, which displays the usage description
<code>-v</code>	Verbose mode. Displays all components, including domain configuration units (DCUs). DCUs include CPU, MCP, hPCI, hsPCI+, and wPCI components. Displays the SC as well.

All of the user privileges groups can use this command; however, the domain administrator and the domain configurator can show boards only in the domains for which they have privileges.

showdevices(1M)

The `showdevices(1M)` command displays the configured physical devices on system boards and the resources that are made available by these devices. Although the `showdevices(1M)` command is not DR-specific, it should be used in conjunction with DR commands. Refer to the section “Showing Device Information” on page 9, and to the `showdevices(1M)` man page, for more information.

The usage information is provided by applications and subsystems that are actively managing system resources. The predicted impact of a system board DR operation can be displayed by performing an offline query of managed resources.

The following table describes the `showdevices(1M)` command options and operands.

Options and Operands	Specifies
<code>board_id</code>	The board ID of the board to be added. The board ID corresponds to the board location (for example, SB2 is the system board in slot 2). Multiple board identifiers are permitted.
<code>-d domain_id domain_tag</code>	Identifies the target domain.
<code>-h</code>	Help, which displays the usage description
<code>-p reports</code>	Offline query information
<code>-v</code>	All I/O devices

Only the domain administrator and the domain configurator can show device information in the domains for which they have privileges.

showplatform(1M)

The `showplatform(1M)` command shows the available component list, the domain state for each domain, and Capacity on Demand (COD) information. Although the `showplatform(1M)` command is not DR-specific, it should be used in conjunction with the DR commands. Refer to the section “Showing Platform Information” on page 11, and to the `showplatform(1M)` man page, for more information.

The following table describes the `showplatform(1M)` command options and operands.

Options and Operands	Specifies
<code>-d domain_id domain_tag</code>	The target domain name
<code>-h</code>	Help, which displays the usage description
<code>-p domains available ethernet cod</code>	Displays reports that are grouped by: domain state (domains); domain available component list (available), domain ethernet addresses (ethernet); or Capacity On Demand (COD) information (COD).
<code>-v</code>	All available command information

All of the user privileges groups, except for the platform service and superuser groups, can use this command; however, the domain administrator and the domain configurator can show the platform information only in the domains for which they have privileges.

Error Message Help System

The SMS software contains an error message help system that you can use to find a description and recovery procedure for a specific error message.

To start the DR error message help system, enter the following command:

```
% /opt/SUNWSMS/jh/scdrhelp &
```

The standard JavaHelp system viewer, `hsviewer`, displays the DR error messages help system. The viewer consists of a toolbar and two panes: the content pane and the navigation pane, as shown in FIGURE 3-1.

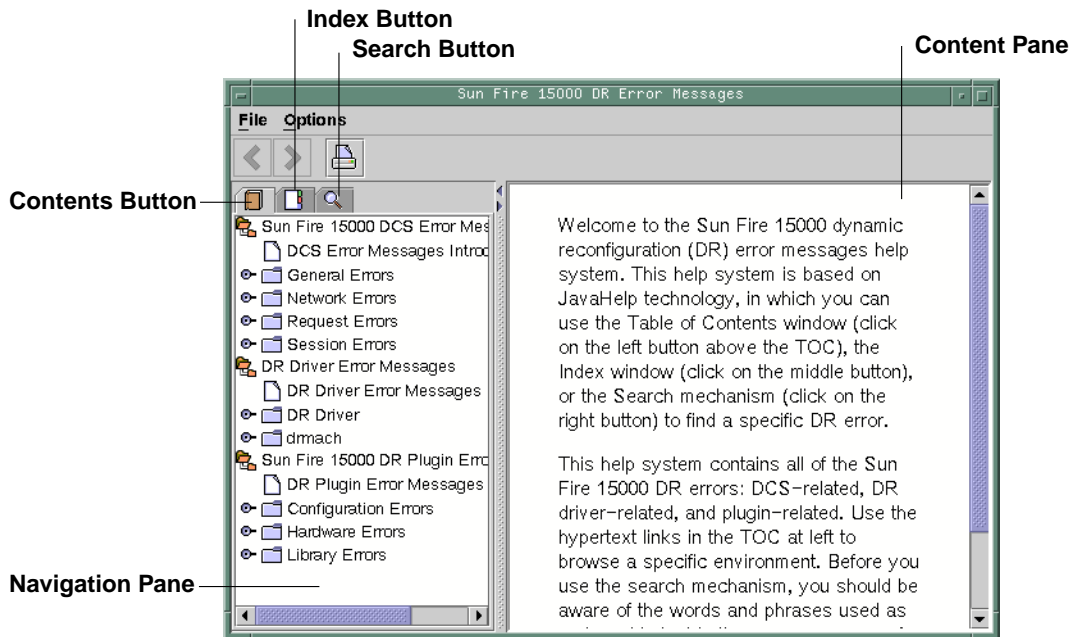


FIGURE 3-1 `hsviewer` GUI Components

JavaHelp Table of Contents

The DR error messages are separated into logical groups according to the type of errors, as shown in FIGURE 3-1. These groups represent the major topics that appear as the top level headings in the table of contents. Error message numbers and/or abbreviated text appear under their respective group name.

JavaHelp Index

The DR error messages are indexed so that key topics are represented in the Index display (FIGURE 3-2). Some index topics are embedded when that embedding is deemed appropriate. For these topics, only the embedded topics are links to error messages.




FIGURE 3-2 JavaHelp Index Display

JavaHelp Search

The DR error messages help system provides a full-text search function. The search database is constructed by indexing the error message help files.

Before you search for a specific error message, first search on a specific string of text in the error message. Also avoid using numeric values because they are treated as replaceable text. The error JavaHelp system window is shown below:

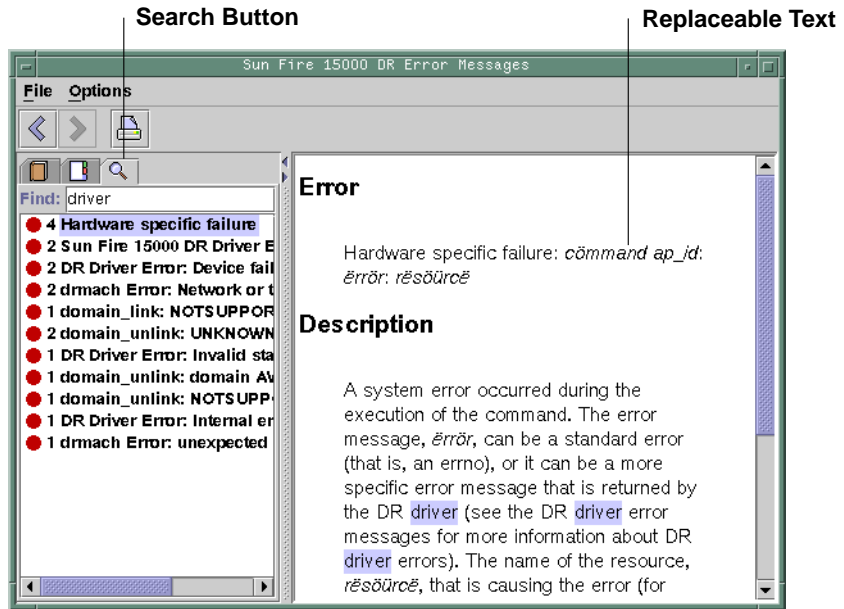


FIGURE 3-3 JavaHelp Search Display

Index

A

- active board state 13
- addboard(1M)
 - example of 21
 - options 20
 - privileges 20
 - procedure 14
- adding boards 14
- administration models 6
- ADR
 - addboard(1M) 19
 - deleteboard(1M) 21
 - moveboard(1M) 22
 - on I/O boards 3
- assigned board state 13
- attachment points 25
 - hardware specific help 25
 - listing 25
- automatic DR 1
- available board state 13
- available component list 5
 - administration models 6
 - board states 13
 - showplatform(1M) output 12

B

- board ID
 - addboard(1M) 20
- board states 12
 - descriptions of 13

- showboards(1M) 26

C

- Capacity on Demand (COD) 4, 14, 16
- cfgadm(1M) 1
- COD 4, 14, 16
- commands
 - addboard(1M) 19
 - cfgadm(1M) 1
 - deleteboard(1M) 21
 - moveboard(1M) 22
 - rcfgadm(1M) 25
 - showboards(1M) 26
 - showdevices(1M) 27
 - showplatform(1M) 26, 28
- components
 - available component list 5
 - types 2
- configuration 21, 23, 25
 - state 20
- configuration state
 - addboard(1M) 20
 - deleteboard(1M) 21
 - moveboard(1M) 23
 - rcfgadm(1M) 25
- CPU component type 2

D

- DCA 6

- deleteboard(1M) 21
 - example 22
 - privileges 22
 - procedure 15
- deleting boards 15
- domain configuration agent (DCA) 6
- domain x server (DXS) 7
- domains
 - assigning slots 5
 - board states 13
 - description of 4
 - logical 5
 - PCD 4
 - physical 5
 - running DR from 1
 - slots 5
- DXS 7
- dynamic reconfiguration 1
- dynamic system domains 4

G
golden IOSRAM 3

H
hardware partitions 4
hardware specific help

- attachment points 25

hot swap 2

I
I/O devices

- component type 2
- reconfiguration of 3

IOSRAM

- golden 3

L
logical domain 5

M
memory component type 2
moveboard(1M)

- description of 22
- example 24
- procedure 16

moving boards 16

O
operands

- addboard(1M) 20
- deleteboard(1M) 21
- moveboard(1M) 23
- rfgadm(1M) 25
- showdevices(1M) 27
- showplatform(1M) 28

options

- addboard(1M) 20
- deleteboard(1M) 21
- moveboard(1M) 23
- rfgadm(1M) 25
- showboards(1M) 27
- showdevices(1M) 27
- showplatform(1M) 28

P
PCD 7
physical domain

- board state 13
- description of 5

platform configuration daemon (PCD) 7
privileges

- model 6
- moveboard(1M) 24
- rfgadm(1M) 26
- showboards(1M) 13, 27
- showdevices(1M) 28
- showplatform(1M) 26, 28

R
rfgadm(1M)

- attachment points 25

- description of 25
- reconfiguration coordination manager 1
- remote DR 25
- replacing system boards 17
- resources 27

S

- showboards(1M)
 - description of 26
 - output 14
 - procedure 12
 - system resources 27
- showdevices(1M)
 - description of 27
 - output 10
 - procedure 9
- showplatform(1M)
 - description of 26, 28
 - output 12
 - procedure 11
- slots 4
- SMS commands 19
- state models 12
- states
 - active 12
 - assigned 12
 - available 12
 - unavailable 12
- system boards
 - replacing 17
 - states 13
- system resources 27

U

- unavailable board state 13

X

- XMITS boards 3

