

Sun Enterprise™ 6x00, 5x00, 4x00, and 3x00 Systems Dynamic Reconfiguration User's Guide



THE NETWORK IS THE COMPUTER™

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
650 960-1300 Fax 650 969-9131

Part No. 806-3984-10
February 2000, Revision A

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

Copyright 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and 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. For Netscape Communicator™, the following notice applies: (c) Copyright 1995 Netscape Communications Corporation. All rights reserved.

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, 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.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

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 2000 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94303-4900 U.S.A. Tous droits réservés.

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. La notice suivante est applicable à Netscape Communicator™: (c) Copyright 1995 Netscape Communications Corporation. Tous droits réservés.

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, 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.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Adobe PostScript

Contents

Preface ix

1. Overview 1

How to Locate Service Procedures and Related Information 1

Sun Enterprise DR Web Site 2

Software Patches 2

Limitations 3

Hardware 3

Hot Plug Support 3

Board Support 3

Broken Boards 4

Non-Detachable Boards 4

Memory Interleaving 4

Permanent Memory 4

Firmware 4

General Support for Dynamic Reconfiguration 4

CPU/Memory Board Firmware 5

Firmware for FC-AL Disk Arrays or Internal Drives 5

Displaying Board Status 5

Basic Status Display 5

Detailed Status Display	6
Terminology	9
The <code>cfgadm</code> Command	9
<code>cfgadm</code> Conditions	10
empty	11
disconnected	11
connected	11
configured	11
unconfigured	12
unknown	12
ok	12
failing	12
failed	13
unusable	13
Naming Conventions for Memory Banks and CPU Numbers	13
Memory Bank <code>ac</code> Numbers	14
CPU Numbers	14
Attachment Point	14
Detachability	15
Conditions and States	16
Connection and Configuration	16
Hot-Plug Hardware	17
Quiescence	17
Suspend-Safe and Suspend-Unsafe Devices	18
Discussion of Board or Device Installation	18
Connecting a Board	19
Configuring a Board	20
Using a Board as a Spare	20

Enabling an Unconfigured Board	20
Addition of Storage Devices	21
Discussion of Board Removal	21
Memory Device Preparation	21
I/O and Network Device Preparation	21
I/O Board Unconfiguration	22
Preparation of an I/O Board for Removal	22
Termination of Network Devices	22
Discussion of Board and Device Replacement or Modification	23
Replacement Sequence	23
Discussion of System Reconfiguration	23
When to Reconfigure	24
I/O Device Reconfiguration	24
Disk Controller Renumbering During a Reconfiguration	24
2. Procedures	27
Displaying PROM Versions	27
Testing for Suspend-Safe Drivers	28
Enabling Dynamic Reconfiguration	28
Removing a Board	29
Removing a CPU/Memory Board	29
Removing an I/O Board	32
Terminating I/O Devices	32
I/O Board Removal	34
Removing Boards That Use Detach-Unsafe Drivers	36
Temporarily Unconfiguring a Board	37
Installing a Board	38
Installing or Replacing a CPU/Memory Board	38

Installing a New I/O Board	41
Installing a Replacement I/O Board	43
Adding Storage Devices	44
Preparing a Spare Board	45
Disabling a Board	45
Enabling Spare Boards	46
Enabling a Single Board	46
Enabling Multiple Boards	47
3. Troubleshooting	49
Troubleshooting Specific Failures	49
Diagnostic Messages	50
Driver Does Not Support Dynamic Reconfiguration	50
Unconfigure Operation Fails	51
CPU/Memory Board Unconfiguration Failure	51
Unable to Unconfigure a Memory Bank	51
Unable to Unconfigure a CPU	52
Unable to Disconnect a Board	53
I/O Board Unconfiguration Failure	53
Device Busy	53
Problems with I/O Devices	54
RPC or TCP Time-out or Loss of Connection	54
Configure Operation Fails	55
CPU/Memory Board Configuration Failure	55
I/O Board Configuration Failure	55
Disabled Board List	56
Glossary	57

Preface

The information in this book is intended for the system administrator and service provider.

This user's guide describes the Dynamic Reconfiguration (DR) feature, which enables you to attach and detach system boards from a running system. The information in this user guide applies to these Sun Enterprise™ systems:

- Sun Enterprise 6500 system
- Sun Enterprise 6000 system
- Sun Enterprise 5500 system
- Sun Enterprise 5000 system
- Sun Enterprise 4500 system
- Sun Enterprise 4000 system
- Sun Enterprise 3500 system
- Sun Enterprise 3000 system

How This Book Is Organized

Chapter 1 gives a general description of Dynamic Reconfiguration (DR).

Chapter 2 provides step-by-step DR procedures.

Chapter 3 has information for troubleshooting DR problems.

Glossary defines the technical terms used in this book.

Typographic Conventions

TABLE P-1 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. Command-line variable; replace with a real name or value.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be <code>root</code> to do this. To delete a file, type <code>rm filename</code> .

Shell Prompts

TABLE P-2 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

TABLE P-3 Related Documentation

Application	Title	Part Number
Sun Enterprise 10000 Dynamic Reconfiguration software	<i>Sun Enterprise 10000 Dynamic Reconfiguration Reference Manual</i>	806-2250-xx
Sun Enterprise 10000 Dynamic Reconfiguration software	<i>Sun Enterprise 10000 Dynamic Reconfiguration User's Guide</i>	806-2249-xx
Sun Enterprise Server Alternate Pathing software	<i>Sun Enterprise Server Alternate Pathing Reference Manual</i>	805-5986-xx
Sun Enterprise Server Alternate Pathing software	<i>Sun Enterprise Server Alternate Pathing User's Guide</i>	805-3532-xx
Sun Management Center 2.1 software	<i>Sun Management Center 2.1 Software User's Guide</i>	806-3166-xx

Accessing Sun Documentation Online

Dynamic Reconfiguration Information

For the latest information about supported hardware, firmware, known bugs, and documentation errata for dynamic reconfiguration, refer to the Solaris™ 8 web page at the web site:

<http://sunsolve2.Sun.COM/sunsolve/Enterprise-dr>

Other Sun Documents

The `docs.sun.com` web site enables you to access Sun technical documentation on the Web. You can browse the `docs.sun.com` archive or search for a specific book title or subject at:

<http://docs.sun.com>

Ordering Sun Documentation

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at:

<http://www1.fatbrain.com/documentation/sun>

Sun Welcomes Your Comments

We are interested in improving our documentation and welcome your comments and suggestions. You can email your comments to us at:

docfeedback@sun.com

Please include the part number of your document in the subject line of your email.

Overview

Dynamic reconfiguration (DR) is an operating environment feature that provides the ability to reconfigure system hardware while the system is running. This feature is optional and can be implemented at the discretion of the system administrator. The main benefit of DR is that a service provider can add or replace hardware resources (such as CPUs, memory, and I/O interfaces) with little interruption of normal system operations.

DR is available for Sun™ system architectures that contain multiple system boards and use board sockets that support hot-plugging. The DR features described in this user's guide are specific to Sun Enterprise™ 6500, 6000, 5500, 5000, 4500, 4000, 3500, and 3000 systems using the Solaris™ 8 operating environment. These features may not apply to other types of server systems.

For information about DR for Sun Enterprise 10000 systems, refer to the *Sun Enterprise 10000 Dynamic Reconfiguration User's Guide* and the *Sun Enterprise 10000 Dynamic Reconfiguration Reference Manual*.

The Sun Management Center™ system monitoring and management software supports dynamic reconfiguration, including features described in this user guide. For more information, refer to the *Sun Management Center 2.1 Software User's Guide*.

Note – For the sake of brevity, the rest of this document refers to an individual system as a “Sun Enterprise xx00 system”, or simply as “the system”.

How to Locate Service Procedures and Related Information

- To determine what types of boards are supported, see “Limitations” on page 3

- To find the system name of a board or device and check its status, see “Displaying Board Status” on page 5
- To install a board, see “Installing a Board” on page 38
- To remove or replace a board, see “Removing a Board” on page 29
- To remove a device driver that does not support Dynamic Reconfiguration, see “Removing Boards That Use Detach-Unsafe Drivers” on page 36
- To connect storage devices to an I/O board, see “Adding Storage Devices” on page 44

Sun Enterprise DR Web Site

For late-breaking news and patch information, visit the Solaris 8 web page at:

<http://sunsolve2.Sun.COM/sunsolve/Enterprise-dr>

The web site is updated periodically.

If you do not have access to this web site, ask your Sun service provider for assistance in obtaining the latest information.

Software Patches

For software patch requirements, visit the Solaris 8 web page at the DR web site noted in the previous section.

Note – SAP R/3 software requires patches to support dynamic reconfiguration. SAP R/3 versions 3.1I and 4.0B currently require the patches `dw1_310.CAR`, `dw2_310.CAR`, and `sapstart`, dated February 1999, but this list is subject to change at any time. Refer to the web page above for any new information about these patches.

Limitations

Hardware

Hot Plug Support

If you see the following message on your console or in your console logs, the hardware cannot be removed while the system is powered up and does not support DR.

```
Hot Plug not supported in this system
```

Board Support

DR may not be fully supported on all board types at this time, although additional support is being developed. For late-breaking news, refer to the Solaris 8 section at the DR web site. See “Sun Enterprise DR Web Site” on page 2.

The `cfgadm` status display may display the following board types, some of which may not be fully supported yet.

TABLE 1-1 Board Types

Type	Name and identifying characteristics
CPU/mem	CPU/memory board with at least one CPU module
Mem	CPU/memory board with no CPU module
Disk board	System board containing a disk drive
Type 1	Dual SBus I/O board with 3 SBus slots
Type 2	SBus-UPA I/O board with 2 SBus slots and 1 frame buffer slot
Type 3	Dual PCI I/O board with 2 PCI card adapter slots
Type 4	SOC+ SBus I/O board with 3 SBus slots
Type 5	SOC+ UPA I/O board with 2 SBus slots, 1 frame buffer slot

Broken Boards



Caution – Inserting a broken (malfunctioning) board may cause a system crash. Use only boards that are known to be functional.

Non-Detachable Boards

If the `cfgadm -v` status display identifies a board as “non-detachable”, the board cannot be dynamically reconfigured. The lowest-numbered CPU/memory board is currently in this category and cannot be removed while the system is running. Support is being developed for these board locations.

Memory Interleaving

Memory boards or CPU/memory boards that contain interleaved memory currently cannot be dynamically reconfigured. To list boards with interleaved memory, use the `prtdiag` or `cfgadm` commands.

Permanent Memory

A CPU/memory board containing non-relocatable memory cannot be dynamically reconfigured. Typically, this condition applies to one CPU/memory board in the system. The board is identified as “PERMANENT” in the status display produced by the `cfgadm -v` command.

Firmware

General Support for Dynamic Reconfiguration

Your machine may require firmware updates to dynamically reconfigure. Look for system messages when the system boots.

Older versions of the CPU PROM may display the following message:

```
Firmware does not support Dynamic Reconfiguration
```

More recent versions of the CPU PROM may display variations of this message.

CPU/Memory Board Firmware

To support DR in the Solaris 8 operating environment, CPU/memory boards may require a PROM upgrade. Instructions for obtaining the CPU upgrade firmware are available at the Solaris 8 section at the DR web site. See “Sun Enterprise DR Web Site” on page 2.

To list board PROM versions, see “Displaying PROM Versions” on page 27.

Firmware for FC-AL Disk Arrays or Internal Drives

For Sun StorEdge™ A5000 disk arrays, or for internal FC-AL disks in the Sun Enterprise 3500 system, the firmware version must be ST19171FC 0413 or later. For more information, refer to the Solaris 8 section at the DR web site. See “Sun Enterprise DR Web Site” on page 2.

Displaying Board Status

The `cfgadm` program displays information about boards and slots. Refer to the `cfgadm(1)` man page for options to this command.

Basic Status Display

Many operations require that you specify the system board names. To obtain these system names, type:

```
# cfgadm
```

When used without options, `cfgadm` displays information about all known attachment points, including memory banks and board slots. The following display shows a typical output.

CODE EXAMPLE 1-1 Output of the Basic `cfgadm` Command

```
# cfgadm
Ap_Id           Receptacle  Occupant    Condition
ac0:bank0       connected   unconfigured ok
ac0:bank1       empty       unconfigured unknown
ac1:bank0       connected   unconfigured ok
```

CODE EXAMPLE 1-1 Output of the Basic `cfgadm` Command (Continued)

```
# cfgadm
ac1:bank1          empty          unconfigured unknown
ac2:bank0          connected      configured   ok
ac2:bank1          empty          unconfigured unknown
ac3:bank0          empty          unconfigured unknown
ac3:bank1          empty          unconfigured unknown
ac4:bank0          empty          unconfigured unknown
ac4:bank1          connected      unconfigured ok
ac8:bank0          empty          unconfigured unknown
ac8:bank1          empty          unconfigured unknown
sysctrl0:slot0    connected      configured   ok
sysctrl0:slot1    connected      configured   ok
sysctrl0:slot2    connected      configured   ok
sysctrl0:slot3    empty          unconfigured unknown
sysctrl0:slot4    empty          unconfigured unusable
sysctrl0:slot5    connected      configured   ok
sysctrl0:slot6    empty          unconfigured unusable
sysctrl0:slot7    empty          unconfigured unknown
sysctrl0:slot8    connected      configured   ok
sysctrl0:slot9    connected      configured   ok
sysctrl0:slot10   connected      configured   ok
sysctrl0:slot11   connected      configured   ok
sysctrl0:slot12   empty          unconfigured unusable
sysctrl0:slot13   disconnected   unconfigured unknown
sysctrl0:slot14   empty          unconfigured unusable
sysctrl0:slot15   disconnected   unconfigured unknown
```

The display lists the memory banks first, followed by information about the board slots. Note that in this example a total of 12 banks are listed, implying there are six CPU/memory boards in the system. (There are two banks of SIMM slots on each Sun Enterprise xx00 CPU/memory board.)

Detailed Status Display

For a more detailed status report, use the command `cfgadm -v`. The `-v` option turns on expanded (verbose) descriptions. CODE EXAMPLE 1-2 is an example of the display produced by the `cfgadm -v` command. Note that example appears to be complicated because the lines wrap around in this display. (This status report is for the same system used in CODE EXAMPLE 1-1.)

CODE EXAMPLE 1-2 Output of the `cfgadm -v` Command

```
# cfgadm -v
```

Ap_Id	Receptacle	Occupant	Condition	Information
When	Type	Busy	Phys_Id	
ac0:bank0	connected	unconfigured	ok	slot0 64Mb base
0xc0000000	disabled-at-boot			
Dec 17 13:30	memory	n	/devices/fhc@0,f8800000/ac@0,1000000:bank0	
ac0:bank1	empty	unconfigured	unknown	slot0 empty
Dec 16 22:42	memory	n	/devices/fhc@0,f8800000/ac@0,1000000:bank1	
ac1:bank0	connected	unconfigured	ok	slot2 1Gb base 0x0
Dec 17 13:30	memory	n	/devices/fhc@4,f8800000/ac@0,1000000:bank0	
ac1:bank1	empty	unconfigured	unknown	slot2 empty
Dec 16 22:42	memory	n	/devices/fhc@4,f8800000/ac@0,1000000:bank1	
ac2:bank0	connected	configured	ok	slot5 1Gb base 0x40000000
permanent				
Dec 16 22:42	memory	n	/devices/fhc@a,f8800000/ac@0,1000000:bank0	
ac2:bank1	empty	unconfigured	unknown	slot5 empty
Dec 16 22:42	memory	n	/devices/fhc@a,f8800000/ac@0,1000000:bank1	
ac3:bank0	empty	unconfigured	unknown	slot8 empty
Dec 16 22:42	memory	n	/devices/fhc@10,f8800000/ac@0,1000000:bank0	
ac3:bank1	empty	unconfigured	unknown	slot8 empty
Dec 16 22:42	memory	n	/devices/fhc@10,f8800000/ac@0,1000000:bank1	
ac4:bank0	empty	unconfigured	unknown	slot11 empty
Dec 16 22:42	memory	n	/devices/fhc@16,f8800000/ac@0,1000000:bank0	
ac4:bank1	connected	unconfigured	ok	slot11 64Mb base
0xc4000000	disabled-at-boot			
Dec 17 13:30	memory	n	/devices/fhc@16,f8800000/ac@0,1000000:bank1	
ac8:bank0	empty	unconfigured	unknown	slot10 empty
Dec 16 22:42	memory	n	/devices/fhc@14,f8800000/ac@0,1000000:bank0	
ac8:bank1	empty	unconfigured	unknown	slot10 empty
Dec 16 22:42	memory	n	/devices/fhc@14,f8800000/ac@0,1000000:bank1	
sysctrl0:slot0	connected	configured	ok	non-detachable
Dec 16 22:42	cpu/mem	n	/devices/central@1f,0/fhc@0,f8800000/clock-board@0,900000:slot0	
sysctrl0:slot1	connected	configured	ok	non-detachable
Dec 16 22:42	dual-sbus	n	/devices/central@1f,0/fhc@0,f8800000/clock-board@0,900000:slot1	
sysctrl0:slot2	connected	configured	ok	
Dec 16 22:42	cpu/mem	n	/devices/central@1f,0/fhc@0,f8800000/clock-board@0,900000:slot2	
sysctrl0:slot3	empty	unconfigured	unknown	
Dec 16 22:42	unknown	n	/devices/central@1f,0/fhc@0,f8800000/clock-board@0,900000:slot3	
sysctrl0:slot4	empty	unconfigured	unusable	
Dec 16 22:42	unknown	n	/devices/central@1f,0/fhc@0,f8800000/clock-board@0,900000:slot4	

CODE EXAMPLE 1-2 Output of the `cfgadm -v` Command

```
sysctrl0:slot5      connected    configured    ok
Dec 16 22:42 cpu/mem      n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot5
sysctrl0:slot6      empty       unconfigured  unusable
Dec 16 22:42 unknown     n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot6
sysctrl0:slot7      empty       unconfigured  unknown
Dec 16 22:42 unknown     n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot7
sysctrl0:slot8      connected    configured    ok
Dec 16 22:42 cpu/mem      n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot8
sysctrl0:slot9      connected    configured    ok
Dec 16 22:42 dual-sbus    n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot9
sysctrl0:slot10     connected    configured    ok
Dec 16 22:42 cpu/mem      n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot10
sysctrl0:slot11     connected    configured    ok
Dec 16 22:42 cpu/mem      n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot11
sysctrl0:slot12     empty       unconfigured  unusable
Dec 16 22:42 unknown     n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot12
sysctrl0:slot13     disconnected unconfigured  unknown      disabled at boot
Dec 16 22:42 dual-sbus    n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot13
sysctrl0:slot14     empty       unconfigured  unusable
Dec 16 22:42 unknown     n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot14
sysctrl0:slot15     disconnected unconfigured  unknown      disabled at boot
Dec 16 22:42 dual-sbus    n          /devices/central@1f,0/fhc@0,f8800000/clock-
board@0,900000:slot15
```

Here are some useful details of the display in FIGURE 1-1:

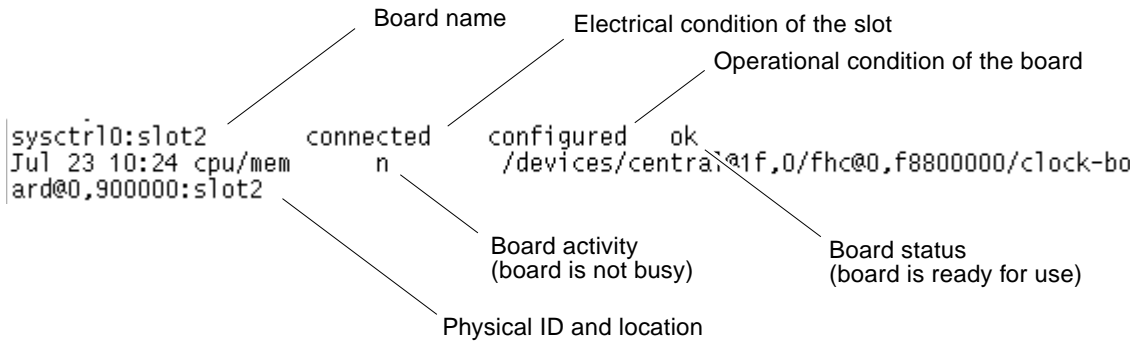


FIGURE 1-1 Details of the Display for `cfgadm -v`

Terminology

The rest of this chapter describes commands and terminology used in DR.

The `cfgadm` Command

In this manual the most frequently used DR command is `cfgadm`.

You can use `cfgadm` to:

- Display board status
- Disable a failing device (remove it from the logical configuration) before the failure can crash the operating environment.
- Add a new or replacement board with minimal interruption to system applications
- Initiate testing of a board
- Change the configuration of boards in the system
- Invoke other hardware-specific functions of a board or related attachment

Many procedures require that you specify the system name for a board. Use the `cfgadm status` report to determine the name and status of the board or card cage slot. For an example, see “Displaying Board Status” on page 5.

The man pages for the `cfgadm` command used on the Sun Enterprise 6x00, 5x00, 4x00, and 3x00 systems include `cfgadm(1M)`, `cfgadm_sysctrl(1M)`, and `cfgadm_ac(1M)`. `cfgadm(1M)` describes the basic functions of the `cfgadm`

command. `cfgadm_sysctrl(1M)` describes additional support for system boards, including newly-added support for CPU/memory boards.
`cfgadm_ac(1M)` describes newly-added support for memory banks.

This release uses a command-line user interface. The Sun Enterprise SyMON system monitoring and management software uses a graphical user interface that supports the DR features described in this user guide. For more information, refer to the *Sun Enterprise SyMON 2.0.1 Software User's Guide*.

Note – DR can work with (but does not require) Alternate Pathing (AP) software. AP switches I/O operations from one I/O board to another. With a combination of DR and AP commands, the system administrator can remove, replace, or deactivate an I/O board with little or no interruption to system operation. Note that for I/O operations, AP requires redundant hardware, meaning that the system must contain an alternate I/O board that is connected to the same device(s) as the board being removed or replaced. For more information on AP, refer to the *Sun Enterprise Server Alternate Pathing User's Guide*.

cfgadm Conditions

The following table lists `cfgadm` conditions for boards and slots. A detailed explanation of each condition and possible corrective actions follow the table.

TABLE 1-2 Summary of Board, Device, and Slot Conditions

Condition	Explanation
empty	No board is present in the slot. All LEDs are off.
disconnected	A board is present but is electrically disconnected.
connected	The board is electrically connected and powered up. The system is actively monitoring the board for temperature and cooling.
configured	Devices on the board are fully initialized and may be mounted or configured for use.
unconfigured	The unconfigured state covers all other device states, including receptacles in the empty state.
unknown	The current condition cannot be determined.
ok	No problems have been detected.

TABLE 1-2 Summary of Board, Device, and Slot Conditions (*Continued*)

Condition	Explanation
failing	A board that was in the OK condition has developed a problem.
failed	The board has failed POST/OBP.
unusable	Either an attachment point has incompatible hardware or an empty attachment point lacks power, cooling, or precharge current.

empty

No board is present in the slot. All LEDs are off.

To install a board, see “Installing a Board” on page 38.

disconnected

A board is present but is electrically disconnected. The system is able to identify the board type. The board LEDs show that the board is in low power mode and can be unplugged at any time.

The LEDs display the following colors: green, yellow, green (Off, On, Off). Use `cfgadm -c disconnect` to enable this state.

To remove a disconnected board, refer to the service manual for the system.

To power up a disconnected board, see “Installing a Board” on page 38.

connected

The board is electrically connected and powered up. The system is actively monitoring the board for temperature and cooling.

The LEDs display the following colors: green, yellow, green (On, Off, Off)

Use `cfgadm -c connect` to enable this state. To remove a connected board, see “Removing a Board” on page 29. To use a connected board, see “Installing a Board” on page 38.

configured

Devices on the board are fully initialized and may be mounted or configured for use. The LEDs show the normal running pattern.

The LEDs display the following colors: On, Off, Flash

Use `cfgadm -c configure` to enable this state.

To remove a configured board, see “Removing a Board” on page 29.

unconfigured

The unconfigured state covers all other device states, including receptacles in the empty state. The LED pattern is the same as for the connected receptacle state.

The LEDs display the following colors: green, yellow, green (On, Off, Off)

Use `cfgadm -c unconfigure` to enable this state.

To remove an unconfigured board, see “Removing a Board” on page 29.

To use an unconfigured board, see “Installing a Board” on page 38.

unknown

The current condition cannot be determined. This situation results either when a new board is inserted in a running system, or a board is placed on the disabled board list prior to a reboot. A transition to a connected receptacle state will change an attachment point condition from unknown to either OK or Failed.

To use an unknown board, see “Installing a Board” on page 38.

ok

No problems have been detected. This condition can only occur after a board has been connected. This condition will persist either until the board is physically removed, or a problem is detected. An ok condition requires correct hardware compatibility, correct firmware revision, adequate power, adequate cooling, and adequate precharge.

To remove an ok board, see “Removing a Board” on page 29.

failing

A failing condition can only occur when a board that was in the OK condition develops a problem. For example, the board has begun to overheat. This condition will be displayed until the problem is corrected or the attachment point is disconnected.

To remove a failing board, see “Removing a Board” on page 29.

To correct an overheating condition, see the system service manual.

failed

The board has failed POST/OBP. A failed condition may occur either during bootup or after a failed connect attempt. This condition is considered uncorrectable and will persist until the board is physically removed. For a failed attachment point condition, the receptacle state should never transition beyond disconnected.

To remove a failed board, see “Removing a Board” on page 29.

unusable

Either an attachment point has incompatible hardware, or an empty attachment point lacks power, cooling, or precharge current. An unusable condition is correctable. This condition is caused by one of the following events:

1. Inadequate cooling in a slot
2. Power is detected in an empty slot
3. A disconnected board has inadequate cooling, inadequate power, or unsupported hardware
4. Firmware has detected a problem either during bootup or when a board is inserted

To remove a board from an unusable slot, see “Removing a Board” on page 29.

To correct overheating conditions in the slot, refer to the system service manual.

Naming Conventions for Memory Banks and CPU Numbers

This section explains the numbering of memory banks and CPUs used in the `cfgadm` status display.

Memory Bank ac Numbers

The `cfgadm` status report lists memory banks in the order of their respective board address controller numbers (`ac0`, `ac1`, `ac2`, and so forth). Note that the ac numbers are not listed in the order of their physical board slot numbers, but in the chronological order in which the CPU/memory boards were inserted into the system. Thus, if the second CPU/memory board is already in slot 7, and you now install a third CPU/memory board in slot 4, a `cfgadm` status report would list the third CPU/memory board (`ac2`) after the second CPU/memory board, even though the third CPU/memory board is in a lower-numbered physical slot.

CPU Numbers

The CPUs are identified by numbers based on the board number. The first CPU number is equal to twice the board number ($2*n$). The second CPU number is twice the board number, plus one ($2*n + 1$).

For example, for board 3 the CPUs are 6 and 7. To see the CPU information for board 3, specify CPUs 6 and 7 in the `psrinfo` command:

```
# psrinfo 6 7
6          on-line   since 01/10/99 18:00:56
7          on-line   since 01/10/99 18:01:01
```

Attachment Point

An *attachment point* is a collective term for a board and its card cage slot.

DR can display the status of the slot, the board, and the attachment point. The DR definition of a board also includes the devices connected to it, so the term *occupant* refers to the combination of board and attached devices.

- A slot (also called a *receptacle*) may have the ability to electrically isolate the occupant from the host machine. That is, the software can put a single slot into low-power mode.
- Receptacles can be named according to slot numbers or can be anonymous (for example, a SCSI chain). To obtain a list of all available logical attachment points, use the `-l` option with the `cfgadm` command.
- An occupant I/O board includes any external storage devices connected by interface cables.

There are two types of system names for attachment points:

- A *physical* attachment point describes the software driver and location of the card cage slot. An example of a physical attachment point name is:

```
/devices/central@1f,0/fhc@0,f8800000/clock-board@0,900000:sysctrl,slot0
```

- A *logical* attachment point is an abbreviated name created by the system to refer to the physical attachment point:

```
sysctrl0:slot0
```

Tip – Note that in the term `sysctrl0`, “1” is a letter and “0” is zero.

Detachability

For a device to be detachable:

- The device driver must support `DDI_DETACH`.
- Critical resources must be redundant or accessible through an alternate pathway. CPUs and memory banks can be redundant critical resources. Disk drives are examples of critical resources that can be accessible through an alternate pathway (through an alternate I/O board).

Some boards cannot be detached. For example, if a system has only one CPU board, that CPU board cannot be detached. An I/O board is not detachable if it controls a boot drive (unless Alternate Pathing is installed on the system, in which case you can switch control of the boot drive to an alternate I/O board).

In the current revision of the software, the lowest-numbered CPU/memory board cannot be detached. In the verbose version of the status display (`cfgadm -v`) these boards are identified as “non-detachable”. For example, in CODE EXAMPLE 1-2, the boards in slot 0 and slot 1 are listed as non-detachable.

If there is no alternate pathway for an I/O board, you can:

- Put the disk chain on a separate I/O board. The secondary I/O board can then be detached.
- Add a second path to the device through a second I/O board. The I/O board can be detached (using Alternate Pathing software to switch access through the alternate board) without losing access to the secondary disk chain.

Conditions and States

A *state* is the operational status of either a receptacle (slot) or an occupant (board).

A *condition* is the operational status of an attachment point.

The `cfgadm` program can display 10 types of states and conditions. See TABLE 1-2.

Note – For a receptacle procedure to be valid, the receptacle must transition in sequence through all three states (empty, disconnected, connected) or in the reverse sequence (connected, disconnected, empty).

Connection and Configuration

There are four main types of operations related to boards:

Connection—in this operation, the slot provides power to the board and begins monitoring the board temperature. For I/O boards, the connection operation is included in the configuration operation (see below). A connection involves a delay that can last up to approximately one minute. The actual time depends on the type of board and the number of boards in the system.

Configuration—the operating environment assigns functional roles to a board and loads device drivers for the board and for devices attached to the board.

Unconfiguration—the system detaches a board logically from the operating environment and takes the associated device drivers offline. Environmental monitoring continues, but any devices on the board are not available for system use.

Disconnection—the system stops monitoring the board and power to the slot is turned off.

If a system board is in use, before powering it off and removing it, stop its use and unconfigure it. After a new or upgraded system board is inserted and powered on, connect its attachment point and configure it for use by the operating environment.

`cfgadm` can connect and configure (or unconfigure and disconnect) in a single command, but if necessary, each operation (connection, configuration, unconfiguration, or disconnection) can be performed separately.

Hot-Plug Hardware

Hot-plug boards and modules have special connectors that supply electrical power to the board or module before the data pins make contact. Boards and devices that do not have hot-plug connectors cannot be inserted or removed while the system is running.

I/O boards and CPU/memory boards used in Enterprise x000 and x500 systems are hot-plug devices. Some devices, such as the clock board and peripheral power supply (PPS), are not hot-plug modules and cannot be removed while the system is running.

Quiescence

During an unconfigure/disconnect operation on a system board with non-pageable OpenBoot™ PROM (OBP) or kernel memory, the operating environment is briefly paused, which is known as operating environment *quiescence*. All operating environment and device activity on the backplane must cease for a few seconds during a critical phase of the operation.

To quiesce a system and test for DR-compatible drivers, see “Testing for Suspend-Safe Drivers” on page 28.

Before it can achieve quiescence, the operating environment must temporarily suspend all processes, CPUs, and device activities. If the operating environment cannot achieve quiescence, it displays the reasons, which may include the following:

- A user thread did not suspend.
- Real-time processes are running.
- A device exists that cannot be paused by the operating environment.

The conditions that cause processes to fail to suspend are generally temporary. Examine the reasons for the failure. If the operating environment encountered a transient condition—a failure to suspend a process—you can try the operation again.

Note – The screen, mouse, and keyboard are not operational while the system is suspended, but you regain control of these devices after the system resumes operation.

Suspend-Safe and Suspend-Unsafe Devices

A *suspend-safe* device is one that does not access memory or interrupt the system while the operating environment is in quiescence. A driver is suspend-safe if it supports operating environment quiescence (suspend/resume). A suspend-safe driver also guarantees that when a suspend request is successfully completed, the device that the driver manages will not attempt to access memory, even if the device is open when the suspend request is made.

Suspend-safe drivers provide the ability to:

- Stop user threads.
- Execute the `DDI_SUSPEND` call in each device driver.
- Stop the clock.
- Stop the CPUs.

A *suspend-unsafe* device allows a memory access or a system interruption while the operating environment is in quiescence.

The operating environment refuses a quiescence request if a suspend-unsafe device is open. To manually suspend the device, you may have to close the device by killing the processes that have it open, by asking users not to use the device, or by disconnecting the cables. For example, if a device that allows asynchronous unsolicited input is open, you can disconnect its cables prior to activating operating environment quiescence and reconnect them after the operating environment resumes. This action prevents traffic from arriving at the device and, thus, the device has no reason to access the backplane.

Tape Devices

The sequential nature of tape devices prevents them from being reliably suspended in the middle of an operation, and then resumed. Therefore, all tape drivers are suspend-unsafe. Before executing an operation that activates operating environment quiescence, make sure all tape devices are closed or not in use.

Discussion of Board or Device Installation

The installation of a new board involves the connection and configuration operations described below. If the board is intended to be a spare board, it must additionally be disabled now, so that you can enable it when you want to use it.

Note – This section does not contain actual procedures. Service procedures begin in Chapter 2.

To install a board, see “Installing a Board” on page 38.

To add a storage device to an existing board, see “Adding Storage Devices” on page 44.

Connecting a Board

After a board is physically inserted into the card cage, a logical connection must be made. For I/O boards the configuration step automatically connects the board. For CPU/memory boards, the connect operation is not included in the configuration step.

The syntax for a board connection is:

```
cfgadm -c connect sysctrl0:slotnumber
```

The term *sysctrl0:slotnumber* is the logical attachment point identification (the system name for the board), which can be found in the `cfgadm` status display.

During the connection process, there is a delay of from 15 seconds to more than a minute before the prompt returns. The length of the delay depends on the type of board and the size and complexity of the system. The system tests the board during this delay.

The states and conditions for the attachment point before a board is inserted are:

- Receptacle state—Empty
- Occupant state—Unconfigured
- Condition—Unknown

After a board is physically inserted, the states and conditions are:

- Receptacle state—Disconnected
- Occupant state—Unconfigured
- Condition—Unknown

After the attachment point is logically connected, the states and conditions are:

- Receptacle state—Connected
- Occupant state—Unconfigured
- Condition—OK

Now the system is aware of the board, but not the usable devices that reside on the board. Temperature is monitored, and power and cooling affect the attachment point condition.

Configuring a Board

For I/O boards the configure operation on a disconnected board will also automatically include the connect operation.

Use the `cfgadm` command to configure a CPU/memory board:

```
# cfgadm -c configure sysctrl0:slotnumber
```

The states and conditions for a configured attachment point are:

- Receptacle state—Connected
- Occupant state—Configured
- Condition—OK

Now the system is also aware of the usable devices which reside on the board and all devices may be mounted or configured for use.

If the configure operation fails for any reason, the states and conditions will still transition to configured. This creates a special situation where the board is partially configured. In this situation, only an *unconfigure* operation is allowed. A further attempt to reconfigure the partial configuration is not permitted.

Using a Board as a Spare

A working board can be kept in the system for use as a spare. To prepare the board for this use, enter the name of the board in the disabled board list. This prevents the board from being used when the system is turned on or rebooted. See “Disabling a Board” on page 45.

To use a spare board, see “Enabling Spare Boards” on page 46.

Enabling an Unconfigured Board

A running system may contain one or more unconfigured boards. That is, the boards are not being used by the system. These unconfigured boards may have been:

- Plugged into the system after the system was booted

- Disabled (as described the previous section)
- Previously unconfigured

To enable a board, use the configure option described in “Configuring a Board” on page 20.

Addition of Storage Devices

To add a storage device, see “Adding Storage Devices” on page 44.

Discussion of Board Removal

The removal of a board requires the devices attached to the board be idled, followed by the unconfiguration and disconnection of the board, as described below.

Note – This section does not contain actual procedures. Service procedures begin in Chapter 2.

The steps include:

1. Preparing the devices on the board.
2. Unconfiguring the board.

Memory Device Preparation

Dynamic reconfiguration of interleaved memory is not currently supported. To determine if interleaved memory is used in the system, use the `prtdiag` or `cfgadm` commands. Memory boards and CPU/memory boards can be dynamically reconfigured if memory on the boards is not interleaved.

I/O and Network Device Preparation

A board with vital system resources cannot be detached unless alternate resources are available on another board. A boot disk is an example of a vital system resource.

A board hosting non-vital system resources can be unconfigured whether or not there are alternate paths to the resources. All its file systems must be unmounted and its swap partitions must be deleted. You may have to kill processes that have open files or devices, or place a hard lock on the file systems (using `lockfs(1M)`) before unmounting the file systems. All I/O device drivers must be detachable.

The system swap space should be configured as multiple partitions on disks attached to controllers hosted by different boards. With this kind of configuration, a particular swap partition is not a vital resource because swap partitions can be added and deleted dynamically. See `swap(1M)` for more information.

Note – When memory or disk swap space is detached, there must be enough memory or swap disk space remaining in the machine to accommodate currently running programs.

I/O Board Unconfiguration

Preparation of an I/O Board for Removal

Before the unconfigure operation can be completed, you must manually terminate usage of all I/O devices on the board, including network interfaces. If Alternate Pathing is installed on your system, switch all I/O functions from the board to alternate I/O boards.

Note – To identify the components that are on the board to be unconfigured, use the `prtdiag(1M)`, `ifconfig(1M)`, `mount(1M)`, `ps(1)`, or `swap(1M)` commands. The `prtdiag(1M)` command provides some information, but is less informative.

Termination of Network Devices

Unconfiguring a board does not automatically terminate use of all network interfaces on the board. You must manually terminate the use of each interface.

You cannot unconfigure any interface that fits the following conditions. In these cases, the unconfigure operation fails with an error message.

- The network interface is the primary network interface for the machine. That is, the IP address of the interface corresponds to the network interface name contained in the file `/etc/nodename`. Halting the primary network interface for the machine prevents network information name services from operating, which

results in the inability to make network connections to remote hosts using applications such as `ftp(1)`, `rsh(1)`, `rcp(1)`, `rlogin(1)`. NFS client and server operations are also affected.

- The interface is the active alternate for an Alternate Pathing (AP) meta device when the AP meta device is plumbed. Interfaces used by the AP system should not be the active path when the board is being unconfigured. Manually switch the active path to one that is not on the board being unconfigured. If no such path exists, manually execute the `ifconfig down` and `ifconfig unplumb` commands on the AP interface. (To manually switch an active path, use the `apconfig(1M)` command.)

Discussion of Board and Device Replacement or Modification

For the procedure to replace a board, see “Installing a Board” on page 38.

For the procedure to add an interface to a board, see “Adding Storage Devices” on page 44.

Replacement Sequence

A number of conditions must be satisfied before a system board can be added to or removed from a system that is under power. For example, the peripheral power supply (PPS) module must be working properly because the PPS supplies precharge current that allows a system board to be safely inserted or removed. A power and cooling module (PCM) must also be working properly in order to supply electrical current and cooling air to system boards.

For these reasons, before you add or replace a system board in Enterprise x000 and x500 servers, first replace any defective PPS or PCM modules.

Discussion of System Reconfiguration

This section discusses reconfiguring your system after you have configured or unconfigured a system board.

When to Reconfigure

In the current version of the software, you might need to reconfigure the system under several conditions, including:

- Board addition—when adding a board, you must execute the reconfiguration sequence to configure the I/O devices associated with the board.
- Board removal—if you remove a board that is not to be replaced, you may (but do not have to) execute the reconfiguration sequence to clean up the `/dev` links for disk devices.
- Board replacement—if you remove a board and then insert it into a different slot, or replace a board with another board that has different I/O devices, you must execute the reconfiguration sequence to configure the I/O devices associated with the board. However, if you replace a board with another board that hosts the *same* set of I/O devices, inserting the replacement into the *same* slot, you may not need to execute the reconfiguration sequence. But be sure to insert a replacement into the same slot that was vacated to retain the original `/dev` link names.

These limitations are expected to be removed in future versions of the software.

I/O Device Reconfiguration

The reconfiguration sequence is the same as the Solaris reconfiguration boot sequence (`boot -r`):

```
drvconfig; devlinks; disks; ports; tapes;
```

When the reconfiguration sequence is executed after a board is configured, device path names not previously seen by the system are entered into the `/etc/path_to_inst` file. The same path names are also added to the `/devices` hierarchy and links to them are created in the `/dev` directory.

Disk Controller Renumbering During a Reconfiguration



Caution – The disk controller number is part of the `/dev` link name used to access the disk. If that number changes during the reconfiguration sequence, the `/dev` link name also changes. This change may affect file system tables and software, such as

Solstice™ DiskSuite™, which uses the `/dev` link names. Update `/etc/vfstab` files and execute other administrative actions necessary due to the changes in the `/dev` link names.

When the reconfiguration sequence is executed after a board is unconfigured or disconnected, the `/dev` links for all the disk partitions on that board may be deleted. The remaining boards retain their current numbering. Disk controllers on a newly inserted board are assigned the next available lowest number by `disks(1M)`.

The `disks(1m)` utility creates symbolic links in the `/dev/dsk` and `/dev/rdsk` directories pointing to the actual special disk device files under the `/devices` directory tree. These entries take the form `/dev/dsk/cxtxdxsx` where:

- `cx` is the disk controller number
- `tx` corresponds to the disk target number, in most cases
- `dx` refers to the logical unit number
- `sx` is the partition number

Removing boards that contain one or more disk controllers prompts the `disks(1m)` utility to examine entries in `/dev/dsk` and `/dev/rdsk`. These entries list the disks attached to the removed controller(s). The `disks(1m)` utility discovers references to disconnected devices have been removed from `/dev/dsk` and `/dev/rdsk`. This removal action makes the logical controller numbers available for re-use. This re-use of controller numbers can lead to confusion when unexpected controller numbers are assigned to disk controllers that are added to the system.

Procedures

These procedures are covered in this chapter:

- “Displaying PROM Versions” on page 27
- “Testing for Suspend-Safe Drivers” on page 28
- “Enabling Dynamic Reconfiguration” on page 28
- “Removing a Board” on page 29
- “Removing Boards That Use Detach-Unsafe Drivers” on page 36
- “Temporarily Unconfiguring a Board” on page 37
- “Installing a Board” on page 38
- “Adding Storage Devices” on page 44
- “Preparing a Spare Board” on page 45

Note – The screen, mouse, and keyboard are not operational at times when DR momentarily suspends the system, but you regain control of these devices when the system resumes operations.

Displaying PROM Versions

- **To see your current PROM version, enter `.version` and banner at the `ok` prompt. Your display may be similar to the following:**

TABLE 2-1 PROM Versions

ok .version							
Board	0:	OBP	3.2.21	199x/06/08	16:58	POST	3.9.4 199x/06/09 16:25
Board	1:	FCODE	1.8.3	199x/11/14	12:41	iPOST	3.4.6 199x/04/16 14:22
Board	2:	FCODE	1.8.7	199x/12/08	15:39	iPOST	3.4.6 199x/04/16 14:22
Board	4:	FCODE	1.8.7	199x/12/08	15:39	iPOST	3.4.6 199x/04/16 14:22

TABLE 2-1 PROM Versions (Continued)

Board 5:	FCODE 1.8.3 199x/11/14 12:41	iPOST 3.4.6 199x/04/16 14:22
Board 6:	FCODE 1.8.7 199x/12/08 15:39	iPOST 3.4.6 199x/04/16 14:22
Board 7:	OBP 3.2.21 199x/06/08 16:58	POST 3.9.4 199x/06/09 16:25

{5} ok **banner**
8-slot Sun Enterprise 4000/5000, No Keyboard
OpenBoot 3.2.21, 1024 MB memory installed, Serial #9039599.
Ethernet address 8:0:xx:xx:xx:xx, Host ID: xxxxxxxx.

Testing for Suspend-Safe Drivers

DR requires board and device drivers that can suspend operations. Such drivers are suspendable or “suspend-safe.”

- **To test for DR-compatible suspendable drivers, use the `quiesce-test` option with the `cfgadm` command:**

```
# cfgadm -x quiesce-test sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “l” is a letter and “0” is zero.

On a large system, the `quiesce-test` command may run as long as a minute or so. During this time no messages are displayed if `cfgadm` does not find incompatible drivers. This is normal behavior.

Enabling Dynamic Reconfiguration

In the `/etc/system` file, two variables must be set to enable dynamic reconfiguration and an additional variable must be set to enable the removal of CPU/memory boards.

1. **Log in as superuser.**

2. To enable dynamic reconfiguration, edit the `/etc/system` file and add the following lines:

```
set pln:pln_enable_detach_suspend=1
set soc:soc_enable_detach_suspend=1
```

3. To enable the removal of a CPU/memory board, edit the `/etc/system` file and add this line:

```
set kernel_cage_enable=1
```

Setting this variable enables the memory unconfiguration operation.

4. Reboot the system to put the changes into effect.

Removing a Board

There are two separate procedures in this section:

- “Removing a CPU/Memory Board” on page 29
- “Removing an I/O Board” on page 32

Removing a CPU/Memory Board

The memory modules on a CPU/memory board can be shared by other CPU/memory boards. It is therefore necessary to halt all use of memory modules on a board before the board can be removed from a system configuration.

Note – The CPU/memory board cannot be removed if (1) it contains interleaved memory or (2) if it is listed in the `cfgadm` status report (`cfgadm -s cols=ap_id:type:info`) as “non-detachable” or “permanent”.

1. Log in as root.

2. Use the `cfgadm` command to determine the system name for the CPU/memory board.

CODE EXAMPLE 2-1 shows the `cfgadm` output for a typical Sun Enterprise 6x00 system.

For the example in this procedure, the board is `ac1`, which has one memory bank (`bank1`).

3. Stop all activity in the memory modules on the board.

This step halts all accesses by other CPU/memory boards and prevents any further use until the board is replaced.

A CPU/memory board can have up to two banks of memory. Memory banks have logical names of the form `acnumber:banknumber`. The term `acnumber` identifies the driver instance, but the number is not directly related to the board slot number. (See “Naming Conventions for Memory Banks and CPU Numbers” on page 13 for an explanation of how the number is derived.) The `banknumber` is either `bank0` or `bank1`. The simple method for determining the names of the memory banks is to examine the output of the following command:

```
# cfgadm -s cols=ap_id:info
```

A typical output is:

TABLE 2-2

Ap_Id	Information
<code>ac0:bank0</code>	<code>slot3 64Mb base 0x0 permanent</code>
<code>ac0:bank1</code>	<code>slot3 empty</code>
<code>ac1:bank0</code>	<code>slot5 empty</code>
<code>ac1:bank1</code>	<code>slot5 64Mb base 0x400000000 disabled-at-boot</code>
<code>sysctrl0:slot1</code>	<code>no ffb installed non-detachable</code>
<code>sysctrl0:slot3</code>	<code>non detachable</code>
<code>sysctrl0:slot5</code>	
<code>sysctrl0:slot7</code>	<code>disabled at boot</code>

This output shows two populated banks of memory: `ac0:bank0` is on the board in `slot3` (`sysctrl0:slot3`) and `ac1:bank` is on the board in `slot 5` (`sysctrl0:slot5`).

In the following example, memory bank 1 is unconfigured on board `ac1`:

```
# cfgadm -c unconfigure ac1:bank1
```

Note – Non-relocatable memory pages in the memory span (a section of memory that is reserved for system use) cannot be unconfigured. Non-relocatable memory is identified as “permanent” in a `cfgadm` listing.

4. To verify that the memory modules are relocatable, use the `cfgadm` command and specify the board name by itself, or the board name and bank number:

```
# cfgadm -v acnumber
```

```
# cfgadm acnumber:banknumber
```

5. Verify that the CPUs on the board are not bound to any processes running in the system.

If a CPU is bound to a process, the board cannot be removed until the process is unbound.

The CPUs are identified by numbers that are related to the board number. The first CPU number is twice the board number ($2*n$). The second CPU number is twice the board number, plus one ($2*n + 1$).

For example, for board 3 the CPUs are 6 and 7. If you wish to see the CPU information for board 3, use the `psrinfo` command and specify CPUs 6 and 7:

```
# psrinfo 6 7  
6          on-line   since 01/10/99 18:00:56  
7          on-line   since 01/10/99 18:01:01
```

To list all bound processes, use the `pbind(1)` command. If any of the listed processes show the CPUs in question, the related boards cannot be removed until those processes are unbound.

6. Unconfigure the board:

```
# cfgadm -c unconfigure sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “l” is a letter and “0” is zero.

7. Disconnect the board:

```
# cfgadm -c disconnect sysctrl0:slotnumber
```

When the LEDs on the board indicate that the board is ready for removal, you can physically remove and replace the board (see “Installing a Replacement I/O Board” on page 43). The two outer LEDs must be off and the middle LED must be on.



Caution – Do not remove a board until it is disconnected, or the system will be damaged.

Tip – If a replacement board is not immediately available, you can leave the board in the system until a replacement arrives.



Caution – If a replacement board is not available and you remove the board, you must fill the empty slot to maintain the proper flow of cooling air in the cardcage. For Sun Enterprise 3000, 3500, 4000, 4500, 5000, and 5500 systems, use a dummy board (part number 504-2592). For Sun Enterprise 6000 or 6500 systems, use a load board (part number 501-3142).

Removing an I/O Board

There are two procedures in this section:

- “Terminating I/O Devices” on page 32
- “I/O Board Removal” on page 34

Terminating I/O Devices

1. If the system is using AP (alternate pathing):
 - a. Switch all board functions to the alternate I/O board.
 - b. Wait until all of the alternate paths are functioning before proceeding.
 - c. Remove the board. See “I/O Board Removal” on page 34.
2. If AP is not available, warn all users to stop using the functions that the board provides.

3. Terminate all usage of devices on the board.

All I/O devices must be closed before they can be unconfigured. Ensure that any networking interfaces on the board are not in use. All storage devices attached to the board should be unmounted and closed. See “I/O Board Unconfiguration” on page 22.

- a. To identify the components that are on the board to be unconfigured, use the `ifconfig`, `mount`, `df`, or `swap` commands.
- b. To see which processes have these devices open, use the `fuser(1M)` command.
- c. Ensure that any networking interfaces on the board are not in use. All storage devices attached to the board should be unmounted and closed.

Note – DR does not automatically terminate network use or close devices. There currently is no way to ensure that the use of the network remains terminated or that all devices remain closed. Other clients may remount them between the time of the unmount and the unconfigure operations.

4. Unmount file systems, including Solstice DiskSuite meta-devices that have a board resident partition, (for example: `umount /partition`)
5. Remove Solstice DiskSuite or Alternate Pathing databases from board-resident partitions. The location of Solstice DiskSuite or Alternate Pathing databases is chosen by the user and can be changed.
6. Remove any private regions used by Sun Enterprise Volume Manager™. The volume manager by default uses a private region on each device that it controls, so such devices must be removed from volume manager control before they can be detached.
7. If the board contains Sun RSM Array™ 2000 controllers, take the controllers offline, using the `rm6` or `rdacutil` commands.
8. Remove disk partitions from the swap configuration.
9. Either kill any process that directly opens a device or raw partition, or direct such a process to close the open device on the board.
10. If a detach-unsafe device is present on the board, close all instances of the device and use `modunload(1M)` to unload the driver. If a detach-unsafe device is present on the board, close all instances of the device and use `modunload(1M)` to unload the driver.



Caution – Unmounting file systems may affect NFS client systems.

I/O Board Removal

1. Terminate all usage of devices on the board.

See “Terminating I/O Devices” on page 32.

2. Check the status of the board:

- For a simple list containing board names, states, and conditions, enter:

```
# cfgadm
```

- For a more detailed list, enter:

```
# cfgadm -v
```

For a board removal or replacement, the states and conditions must be one of the following sets:

- The board is ok:
 - Receptacle state—Connected
 - Occupant state—Configured
 - Condition—OK
- The board is failing:
 - Receptacle state—Connected
 - Occupant state—Configured
 - Condition—Failing

3. Unconfigure the board:

```
# cfgadm -c unconfigure sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “1” is a letter and “0” is zero.

For `sysctrl0:slotnumber` (the attachment point ID) use the board name that was listed in the status report of the previous step.

For an I/O board, the unconfigure operation normally also disconnects the board.

4. Use the `cfgadm` command to confirm that the board is unconfigured.

If the unconfigure operation failed:

- a. See “Removing Boards That Use Detach-Unsafe Drivers” on page 36.

- b. See “Quiescence” on page 17.
- c. Resolve the problem.
- d. Unconfigure the board again (Step 1).

Note – A failure of the unconfigure step results in a partially unconfigured condition. If this happens, attempt to unconfigure again. A configuration operation is not permitted at this point.

5. When the board is unconfigured, you can do one of the following:

- Leave the board in the system unconfigured.
- Configure the board.
- Disconnect the board manually, if the unconfiguration operation did not do so automatically:

```
# cfgadm -v -c disconnect sysctrl0:slotnumber
```

6. If you wish to remove the board from the card cage, first verify the board status.

- a. Use `cfgadm` to verify that the board is logically disconnected.
- b. Check the LEDs on the board to verify that the board is electrically disconnected.

The two outer LEDs must be off and the middle LED must be on.

After you have verified that the board is disconnected, and the peripheral power supply is operating properly (see “Replacement Sequence” on page 23), you can physically remove or replace the board. For the replacement procedure, see “Installing a Board” on page 38.

If a replacement board is not available, you can leave the board in the system until a replacement arrives.



Caution – If you remove a board and a replacement board is not immediately available, you must fill the empty slot to maintain the proper flow of cooling air in the cardcage. For Sun Enterprise 3000, 3500, 4000, 4500, 5000, and 5500 systems, use a dummy board (part number 504-2592). For Sun Enterprise 6000 or 6500 systems, use a load board (part number 501-3142).

Removing Boards That Use Detach-Unsafe Drivers

Some drivers do not yet support DR on Sun Enterprise 3x00, 4x00, 5x00, and 6x00 systems. DR cannot detach these drivers, but you can remove some undetachable drivers manually.

1. **Halt all use of the device controller.**
2. **Halt the use of all other controllers of the same type on all boards in the machine.**
The remaining controllers can be used again after the DR unconfigure operation is complete.
3. **Use appropriate Unix commands to manually close all such drivers on the board.**
4. **Use the `modinfo(1M)` command to find the module IDs of the drivers, then use the `modunload(1M)` command to unload them.**
5. **Disconnect the board with this command:**

```
# cfgadm -c disconnect sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “l” is a letter and “0” is zero.

The disconnected board can be physically removed now or at a later time.



Caution – If you remove a board and a replacement board is not immediately available, you must fill the empty slot to maintain the proper flow of cooling air in the cardage. For Sun Enterprise 3000, 3500, 4000, 4500, 5000, and 5500 systems, use a dummy board (part number 504-2592). For Sun Enterprise 6000 or 6500 systems, use a load board (part number 501-3142).

Tip – If you cannot execute the above steps, recover the system configuration by adding the board to the disabled board list using the NVRAM setting `disabled-board-list` (see Platform Notes), then reboot the system. Remove the board at a later time.

Tip – Many third-party drivers (those purchased from vendors other than Sun Microsystems) do not yet properly support the standard Solaris software `modunload` interface. Test these driver functions during the qualification and installation phases of any third-party device.

Temporarily Unconfiguring a Board

If a replacement board or a filler board (a dummy board or a load board, where applicable) is not available, you can use DR to power down the board and leave it in place.

1. **Prepare the board with the procedures in “Discussion of Board Removal” on page 21.**

Note – To identify the components that are on the board to be unconfigured, use the `ifconfig`, `mount`, `df`, or `swap` commands. Another somewhat less informative way is to execute the `prtdiag(1M)` command.

2. **Make sure the device is not being used.**

For a board removal or replacement, the states and conditions must be one of the following sets:

- The board is ok:
 - Receptacle state—Connected
 - Occupant state—Configured
 - Condition—OK
- The board is failing:
 - Receptacle state—Connected
 - Occupant state—Configured
 - Condition—Failing

3. **Unconfigure the attachment point occupant:**

```
# cfgadm -v -c unconfigure sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “l” is a letter and “0” is zero.

Note – If the unconfigure step fails, the states and conditions will remain the same as before. This creates a special situation in which the board is only partially unconfigured. In this situation, attempt to unconfigure again. An attempt to configure or reconfigure is not permitted at this point.

4. Disconnect the attachment point:

```
# cfgadm -v -c disconnect sysctrl0:slotnumber
```

5. If you do not want the attachment point to be enabled at boot:

```
# cfgadm -o disable-at-boot sysctrl0:slotnumber
```

Installing a Board

When installing a board:

- Do not use a board that is bad or suspected to be unreliable. It can crash the system.
- The board PROM version must support DR functionality.
- The board type and option cards must be supported by DR. Refer to the web site for the current list of supported hardware.

There are three separate procedures in this section;

- “Installing or Replacing a CPU/Memory Board”
- “Installing a New I/O Board” on page 41
- “Installing a Replacement I/O Board” on page 43

Installing or Replacing a CPU/Memory Board

- 1. If the peripheral power supply (PPS) is faulty, replace it before beginning this procedure. The PPS must be able to supply precharge current to the board that is being installed or removed.**

2. Verify that the selected board slot can accept a board.

```
# cfgadm
```

The states and conditions should be:

- Receptacle state—Empty
- Occupant state—Unconfigured
- Condition—Unknown

or

- Receptacle state—Disconnected
- Occupant state—Unconfigured
- Condition—Unknown

3. Physically insert the board into the slot and watch for an acknowledgment on the system console or in the system log file. The acknowledgment is of the form “*name* board inserted into slot3”.

After a CPU/memory board is inserted, the states and conditions should become:

- Receptacle state—Disconnected
- Occupant state—Unconfigured
- Condition—Unknown

Any other states or conditions are an error.

4. Configure the board:

```
# cfgadm -v -c configure sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “1” is a letter and “0” is zero.

There is a delay of about a minute before the message appears. The system is testing the board during the delay.

The states and conditions for a connected and configured attachment point should be:

- Receptacle state—Connected
- Occupant state—Configured
- Condition—OK

Now the system is aware of the usable devices on the board and the devices can be used.

5. Configure the memory devices on the board:

```
# drvconfig -i ac
```

6. Determine the system numbers of the new CPU modules. For example:

CODE EXAMPLE 2-1 Using `psrinfo` to List CPU Module System Numbers

```
# psrinfo
6      on-line   since 12/08/98 11:01:25
7      on-line   since 12/08/98 11:01:29
10     powered-off since 12/08/98 12:42:17
```

In this example, there is one new CPU module (system number 10). The module has not yet been enabled, so it is listed as being powered off.

Note – The system number for a CPU is calculated from the board number and is equal to twice the board number, plus 0 for CPU module 0, or 1 for CPU module 1. In CODE EXAMPLE 2-1 system number 10 represents module 0 on board number 5.

7. Enable the new CPU module or modules:

```
# psradm -n number number
```

In CODE EXAMPLE 2-1, there is only one CPU module (10), so the command is:

```
# psradm -n 10
```

8. Test the new memory banks:

```
# cfgadm -o test_type -t acnumber:bank0
# cfgadm -o test_type -t acnumber:bank1
```

where *test_type* is one of three memory tests:

- Quick (writes a pattern of ones and zeros)
- Normal (detects specific memory address failures)
- Extended (tests interference between memory cells)

Note – For one Gbyte of memory the test times are on the order of several minutes (for the quick and normal tests) to more than six hours (for the extended test).

To determine the logical names of the new board, see Step 1 in “Removing a CPU/Memory Board” on page 29.

9. Configure the new memory banks:

```
# cfgadm -c configure acnumber:bank0
# cfgadm -c configure acnumber:bank1
```

10. Verify that the board and the memory banks are configured.

- For the CPU status, use the `psrinfo` or `mpstat` commands.
- For the memory status, use the `prtconf` or `vmstat` commands.

Installing a New I/O Board

1. If the peripheral power supply (PPS) is faulty, replace it before beginning this procedure. The PPS must be able to supply precharge current to the board that is being installed or removed.
2. Verify that the selected board slot is ready for a board.

```
# cfgadm
```

The states and conditions should be:

- Receptacle state—Empty
- Occupant state—Unconfigured
- Condition—Unknown

or

- Receptacle state—Disconnected
- Occupant state—Unconfigured
- Condition—Unknown

3. Physically insert the board into the slot and look for an acknowledgment on the console, such as, “*name* board inserted into slot3”.

After an I/O board is inserted, the states and conditions should become:

- Receptacle state—Disconnected
- Occupant state—Unconfigured

- Condition—Unknown

Any other states or conditions should be considered an error.

4. Connect any peripheral cables and interface modules to the board.

5. Configure the board with the command:

```
# cfgadm -v -c configure sysctrl0:slotnumber
```

Tip – In the term *sysctrl0*, “l” is a letter and “0” is zero.

This command should both connect and configure the receptacle. Verify with the *cfgadm* command.

The states and conditions for a connected and configured attachment point should be:

- Receptacle state—Connected
- Occupant state—Configured
- Condition—OK

Now the system is also aware of the usable devices which reside on the board and all devices may be mounted or configured to be used.

If the command fails to connect and configure the board and slot (the status should be shown as “configured” and “ok”), do the connection and configuration as separate steps:

a. Connect the board and slot by entering:

```
# cfgadm -v -c connect sysctrl0:slotnumber
```

There is a delay of 15 seconds or more before the message appears. The system is testing the board during the delay.

The states and conditions for a connected attachment point should be:

- Receptacle state—Connected
- Occupant state—Unconfigured
- Condition—OK

Now the system is aware of the board, but not the usable devices which reside on the board. Temperature is monitored and power and cooling affect the attachment point condition.

b. Configure the board and slot by entering:

```
# cfgadm -v -c configure sysctrl10:slotnumber
```

The states and conditions for a configured attachment point should be:

- Receptacle state—Connected
- Occupant state—Configured
- Condition—OK

Now the system is also aware of the usable devices which reside on the board and all devices may be mounted or configured to be used.

6. Reconfigure the devices on the board by entering:

```
# drvconfig; devlinks; disks; ports; tapes;
```

The console should display a list of devices and their addresses.

7. Activate the devices on the board using commands such as `mount` and `ifconfig`, as appropriate.

Installing a Replacement I/O Board

1. If you are not continuing from “Removing an I/O Board” above, use the `cfgadm` command and select a card cage slot to use, but do not insert the board yet.
2. View the configuration list and verify that the slot is unconfigured:

```
# cfgadm
```

3. Insert the board in the slot and look for an acknowledgment on the console, such as, “*name* board inserted into slot3.”
4. Use the `cfgadm` command again to look for the system name assigned to the new board.
5. Configure the board using the system name for the board:

```
# cfgadm -c configure sysctrl10:slotnumber
```

Tip – In the term `sysctrl0`, “1” is a letter and “0” is zero.

There is a delay of 15 seconds or more before the message appears. The system is testing the board during the delay.

6. **Configure any I/O devices on the board using commands such as `drvconfig` and `devlinks`, as appropriate.**
7. **Activate the devices on the board using commands such as `mount` and `ifconfig`, as appropriate.**

Adding Storage Devices

To add storage devices to an I/O board in the system:

1. **Terminate all active use of the devices on the I/O board.**
2. **Unconfigure the board.**

```
# cfgadm -c unconfigure sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “1” is a letter and “0” is zero.

3. **To remove the board from the card cage:**
 - a. **Use `cfgadm` to verify that the board is logically disconnected.**
 - b. **Check the LEDs on the board to verify that the board is electrically disconnected.**

The two outer LEDs must be off and the middle LED must be on.
 - c. **Physically remove the board.**
4. **Add the storage device controller:**
 - For an optical controller, attach the I/O module and interface cable.
 - For an SBus or PCI controller card, use the Disconnect command before removing the board. Add the controller card and place the I/O board back in the card cage.

5. **Insert the board into the slot and watch for an acknowledgment on the system console or in the system log file. The acknowledgment is of the form, “*name* board inserted into slot3.”**

After a CPU/memory board is inserted, the states and conditions should become:

- Receptacle state—Disconnected
- Occupant state—Unconfigured
- Condition—Unknown

Any other states or conditions should be considered an error.

6. **Reconfigure the board.**

```
# cfgadm -c configure sysctrl0:slotnumber
```

There is a delay of 15 seconds or more before the prompt reappears. The system is testing the board during the delay.

Only the Occupant state should change. The Receptacle state and condition should remain the same.

7. **If you installed the board in a different slot, reconfigure the devices on the board by entering:**

```
# drvconfig; devlinks; disks; ports; tapes;
```

The console should display a list of devices and their addresses.

8. **Activate the devices on the board using commands such as `mount` and `ifconfig`, as appropriate.**

Preparing a Spare Board

A working board can be disabled for later use as a spare.

Disabling a Board

There are two methods for disabling a board. You can use an EEPROM command or a `cfgadm` command.

- To use the EEPROM command to disable a board:

```
# eeprom disabled-board-list=sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “1” is a letter and “0” is zero.

- To use the `cfgadm` command to disable a board:

```
# cfgadm -c disconnect -o disable-at-boot sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “1” is a letter and “0” is zero.

Enabling Spare Boards

Enabling a Single Board

You can enable a single disabled board immediately, or set it to be enabled at the next boot:

- **To immediately override the disabled condition, use the force flag (`-f`) with the `cfgadm` command:**

```
# cfgadm -f -c connect sysctrl0:slotnumber
```

- **To enable at boot, use the enable option (`-o enable-at-boot`) with the `cfgadm` command:**

```
# cfgadm -o enable-at-boot -c connect sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “1” is a letter and “0” is zero.

Enabling Multiple Boards

You can set all boards to be enabled at the next boot:

- **If you are at the system prompt, use the `eeeprom` command to remove all boards from the disabled board list by setting the `disabled-board-list` variable to a null set:**

```
# eeeprom disabled-board-list=
```

- **If you are at the OpenBoot prompt, use this OBP command to remove all boards from the disabled board list:**

```
OK set-default disabled-board-list
```


Troubleshooting

Troubleshooting Specific Failures

This chapter discusses common types of failure:

- “Driver Does Not Support Dynamic Reconfiguration” on page 50
- “Unconfigure Operation Fails” on page 51
- “Configure Operation Fails” on page 55

Diagnostic Messages

The following are examples of `cfgadm` diagnostic messages. (Syntax error messages are not included here.)

```
cfgadm: Configuration administration not supported on this machine
cfgadm: hardware component is busy, try again
cfgadm: operation: configuration operation not supported on this machine
cfgadm: operation: Data error: error_text
cfgadm: operation: Hardware specific failure: error_text
cfgadm: operation: Insufficient privileges
cfgadm: operation: Operation requires a service interruption
cfgadm: System is busy, try again
cfgadm: Hardware specific failure: memory delete failed: VM viability test
failed
cfgadm: Hardware specific failure: memory delete failed: memory operation
refused
cfgadm: Hardware specific failure: memory delete failed: memory delete timeout
WARNING: Processor number number failed to offline.
NOTICE: dual-sbus-soc+ board in slot 4 partially configured
```

See `config_admin(3X)` for additional error message detail.

Driver Does Not Support Dynamic Reconfiguration

Some drivers do not yet support quiesce operations. A DR-compatible driver must be suspendable.

- Use this command to test for suspendable drivers.

```
# cfgadm -x quiesce-test sysctrl0:slotnumber
```

Tip – In the term `sysctrl0`, “1” is a letter and “0” is zero.

DR may not yet support some types of I/O boards in Sun Enterprise 6x00, 5x00, 4x00, and 3x00 systems. For late-breaking news, refer to the Solaris 8 section at the DR web site. See “Sun Enterprise DR Web Site” on page 2.

Unconfigure Operation Fails

An unconfigure operation can fail if:

- The devices on the board are in use.
- The affected drivers are not detachable.

CPU/Memory Board Unconfiguration Failure

Problems that prevent unconfiguration for the CPU/memory board are:

- The memory banks on the board are configured (in use). See “Unable to Unconfigure a Memory Bank” on page 51.
- CPUs on the board cannot be taken off line. See “Unable to Unconfigure a CPU” on page 52.
- The board cannot be disconnected after it is unconfigured. See “Unable to Disconnect a Board” on page 53.

Unable to Unconfigure a Memory Bank

To unconfigure a memory bank, it must be possible to move the contents of the memory to the swap device, file system or some other piece of memory that is not being deleted.

Bank Cannot Be Reconfigured

If the unconfigure fails with the following message, this bank cannot be unconfigured:

```
cfgadm: Hardware specific failure: memory delete failed: non-relocatable pages
in span
```

Some memory pages that cannot be moved.

- **To confirm that a memory page cannot be moved, use the verbose option with the `cfgadm` command and look for the word “permanent” in the listing:**

```
# cfgadm -v acnumber
```

Not Enough Available Memory

If the unconfigure fails with one of the messages below, there would not enough available memory in the system if the board is removed:

```
cfgadm: Hardware specific failure: memory delete failed: VM viability test failed
```

```
cfgadm: Hardware specific failure: memory delete failed: memory operation refused
```

- **Reduce the memory load on the system and try again. If practical, install more memory in another board slot.**

Memory Demand Increased

If the unconfigure fails with the following message, the memory demand has increased while the unconfigure operation was proceeding:

```
cfgadm: Hardware specific failure: memory delete failed: memory delete timeout
```

- **Reduce the memory load on the system and try again.**

Unable to Unconfigure a CPU

CPU unconfiguration is part of the unconfiguration operation for a CPU/memory board. If the operation fails to take the CPUs offline, the following message is logged to the console:

```
WARNING: Processor number failed to offline.
```

This failure occurs if:

- The CPU has processes bound to it.
- The CPU is the last one in a CPU set.
- The CPU is the last on-line CPU in the system.

Unable to Disconnect a Board

It is possible to unconfigure a board and then discover that it cannot be disconnected. The `cfgadm` status display lists the board as not detachable. This problem occurs when the board is supplying an essential hardware service that cannot be relocated to an alternate board.

I/O Board Unconfiguration Failure

A device cannot be unconfigured or disconnected while it is in use. Many failures to unconfigure I/O boards occur because activity on the boards have not been stopped, or because an I/O device becomes active again after it has been stopped.

If Alternate Pathing is in use on the system, switch all I/O activity from the board to the alternate I/O board.

Device Busy

Disks attached to an I/O board must idled before any attempt is made to unconfigure or disconnect that board. Any attempt to unconfigure/disconnect a board whose devices are still in use will be rejected.

If an unconfiguration operation fails because an I/O board has a busy or open device, the board is left only partially unconfigured. The operation sequence stopped at the busy device.

To regain access to the devices which were not unconfigured, the board must be completely unconfigured and then reconfigured.

In such a case, the system will log messages similar to the following:

```
NOTICE: unconfiguring dual-pci board in slot 7
NOTICE: dual-pci board in slot 7 partially unconfigured
```

- **To continue the unconfigure operation, unmount the device and retry the unconfigure operation. The board must be in the unconfigured state before you try to reconfigure this board.**

Problems with I/O Devices

All I/O devices must be closed before they are unconfigured.

1. **To see which processes have these devices open, use the `fuser(1M)` command.**
2. **Perform the following tasks for I/O devices.**
 - If the redundancy features of Alternate Pathing or Solstice DiskSuite mirroring are used to access a device connected to the board, reconfigure these subsystems so that the device or network is accessible by way of controllers on other system boards.
 - Unmount file systems, including Solstice DiskSuite meta-devices that have a board resident partition. (For example, `umount /partition`).
 - Remove Solstice DiskSuite or Alternate Pathing databases from board-resident partitions. The location of Solstice DiskSuite or Alternate Pathing databases is explicitly chosen by the user and can be changed.
 - Remove any private regions used by Sun Volume Manager or Veritas Volume Manager. Volume Manager by default uses a private region on each device that it controls, so such devices must be removed from Volume Manager control before they can be detached.
 - Take offline any RSM 2000 controllers on the board that is being detached, using the `rm6` or `rdacutil` commands.
 - Remove disk partitions from the swap configuration.
 - Either kill any process that directly opens a device or raw partition, or direct it to close the open device on the board.
 - If a detach-unsafe device is present on the board, close all instances of the device and use `modunload(1M)` to unload the driver. If a detach-unsafe device is present on the board, close all instances of the device and use `modunload(1M)` to unload the driver.



Caution – Unmounting file systems may affect NFS client systems.

RPC or TCP Time-out or Loss of Connection

Time-outs occur by default after two minutes. Administrators may need to increase this time-out value to avoid time-outs during a DR-induced operating system quiescence, which may take longer than two minutes. Quiescing a system makes the system and related network services unavailable for a period of time that can exceed two minutes. These changes affect both the client and server machines.

Configure Operation Fails

CPU/Memory Board Configuration Failure

The attempt to configure a memory bank fails if the board has been intentionally disabled. For example:

```
# cfgadm -c configure ac0:bank0
cfgadm: Hardware specific failure: memory is disabled at boot
```

- Use the `-f` (force) option to overcome this problem.

```
# cfgadm -c configure -f ac0:bank0
```

I/O Board Configuration Failure

A configure operation may fail because an I/O board with a device does not currently support hot-plugging. In such a situation, the board is now only partially configured. The operation has stopped at the unsupported device. In this situation, the board must be brought back to the unconfigured state before another configure attempt. In such a case the system will log messages similar to the following:

```
NOTICE: configuring dual-sbus-soc+ board in slot 4
NOTICE: dual-sbus-soc+ board in slot 4 partially configured
```

- To continue the configure operation, either remove the unsupported device's driver or replace it with a new version of the driver that will support hot-plugging.

Disabled Board List

Attempting to connect a board may produce the following error message, if the board is on the disabled board list:

```
# cfgadm -c connect sysctrl0:slotnumber
cfgadm: Hardware specific failure: connect failed: board is
disabled: must override with [-f][--o enable-at-boot]
```

- **To override the disabled condition, use the force flag (-f) or the enable option (--o enable-at-boot) with the `cfgadm` command, as shown below:**

```
# cfgadm -f -c connect sysctrl0:slotnumber
```

```
# cfgadm --o enable-at-boot -c connect sysctrl0:slotnumber
```

- **To remove all boards from the disabled board list, set the `disabled-board-list` variable to a null set by entering the system command:**

```
# eeprom disabled-board-list=
```

- **If you are at the OpenBoot prompt, use this OBP command instead to remove all boards from the disabled board list:**

```
OK set-default disabled-board-list
```

Glossary

AP See Alternate Pathing.

ac Address controller. The `cfgadm` status report lists memory banks in the order of the board address controller numbers (`ac0`, `ac1`, `ac2`, and so forth). Note that the `ac` numbers are not listed in the order of their physical board slot numbers, but in the chronological order in which the CPU/memory boards were inserted into the system. Thus, if the second CPU/memory board is already in slot 7, and you now install a third CPU/memory board in slot 4, a `cfgadm` status report would list the third CPU/memory board (`ac2`) after the second CPU/memory board, even though the third CPU/memory board is in a lower-numbered physical slot.

ap_id Attachment point identifier; an `ap_id` specifies the type and location of the attachment point in the system and is unambiguous. There are two types of identifiers: physical and logical. A physical identifier contains a fully specified pathname, while a logical identifier contains a shorthand notation.

Alternate Pathing Alternate Pathing (AP) is software package that allows the use of multiple paths between a server and a disk array or a network. If one path fails, AP can ensure that the disk array or network is still available through the alternate path. For example, the alternate path can be a second port on an interface board, or an entirely separate interface board. See also Dynamic Reconfiguration.

Attachment point A collective term for a board and its card cage slot. A *physical* attachment point describes the software driver and location of the card cage slot. A *logical* attachment point is an abbreviated name created by the system to refer to the physical attachment point.

cfgadm command `cfgadm` is the primary command for dynamic reconfiguration on the Sun Enterprise 6x00, 5x00, 4x00, and 3x00 systems. For information about the command and its options, refer to the `cfgadm(1M)`, `cfgadm_sysctrl(1M)`, and `cfgadm_ac(1M)` man pages. For any late-breaking news about this and related commands, refer to the Solaris 8 section at the DR web site. See “Sun

Enterprise DR Web Site” on page 1-2. For dynamic reconfiguration commands used on the Sun Enterprise 10000 system, refer to the *Sun Enterprise 10000 Dynamic Reconfiguration User’s Guide*.

Condition The operational status of an attachment point.

Configuration (system) The collection of attached devices known to the system. The system cannot use a physical device until the configuration is updated. The operating system assigns functional roles to a board and loads device drivers for the board and for devices attached to the board.

Configuration (board) The operating system assigns functional roles to a board and loads device drivers for the board and for devices attached to the board.

Connection A board is present in a slot and is electrically connected. The temperature of the slot is monitored by the system.

Detachability The device driver supports DDI_DETACH and the device (such as an I/O board or a SCSI chain) is physically arranged so that it can be detached.

Disconnection The system stops monitoring the board and power to the slot is turned off. A board in this state can be unplugged.

DR See Dynamic Reconfiguration

Dynamic Reconfiguration Dynamic Reconfiguration (DR) is software that allows the administrator to (1) view a system configuration; (2) suspend or restart operations involving a port, storage device, or board; and (3) reconfigure the system (detach or attach hot-swappable devices such as disk drives or interface boards) without the need to power down the system. When DR is used with Alternate Pathing or Solstice DiskSuite software (and redundant hardware), the server can continue to communicate with disk drives and networks without interruption while a service provider replaces an existing device or installs a new device. DR supports replacement of a CPU/Memory, provided the memory on the board is not interleaved with memory on other boards in the system.

Hot-plug Hot-plug boards and modules have special connectors that supply electrical power to the board or module before the data pins make contact. Boards and devices that do not have hot-plug connectors cannot be inserted or removed while the system is running.

Hot swap A hot swap device has special DC power connectors and logic circuitry that allow the device to be inserted without the necessity of turning off the system.

Logical DR	A DR operation in which hardware is not physically added or removed. An example is the deactivation of a failed board that is then left in the slot (to avoid changing the flow of cooling air) until a replacement is available.
Physical DR	A DR operation that involves the physical addition or removal of a board. See also “Logical DR.”
Quiescence	A brief pause in the operating environment to allow an unconfigure and disconnect operation on a system board with non-pageable OpenBoot PROM (OBP) or kernel memory. All operating environment and device activity on the backplane must cease for a few seconds during a critical phase of the operation.
Receptacle	A receiver such as a board slot or SCSI chain.
State	The operational status of either a receptacle (slot) or an occupant (board).
Suspendability	To be suitable for DR, a device driver must have the ability to stop user threads, execute the <code>DDI_SUSPEND</code> call, stop the clock, and stop the CPUs.
Suspend-safe	A suspend-safe device is one that does not access memory or interrupt the system while the operating system is in quiescence. A driver is considered suspend-safe if it supports operating system quiescence (suspend/resume). It also guarantees that when a suspend request is successfully completed, the device that the driver manages will not attempt to access memory, even if the device is open when the suspend request is made.
Suspend-unsafe	A suspend-unsafe device is one that allows a memory access or a system interruption while the operating system is in quiescence.
SyMON	Sun Enterprise SyMON is a graphical user interface for monitoring and managing systems. The interface includes dynamic reconfiguration capability.
Occupant	Hardware resource such as a system board or a disk drive that occupies a DR receptacle or slot.
Unconfiguration	The system detaches a board logically from the operating system and takes the associated device drivers off-line. Environmental monitoring continues, but any devices on the board are not available for system use.

Index

B

board, preparation for detaching, 22

C

cfgadm

cfgadm -c configure, 19, 20, 31, 32, 36, 39, 40, 42, 43, 45

cfgadm -c connect, 19

cfgadm -v, 6

cfgadm -v -c connect, 42

cfgadm -v -c disconnect, 35, 38

cfgadm -v -l, for expanded hardware information, 34

D

detach command and non-network devices, 54

detaching a board, preparation for, 22

detaching network device, preparation for, 22

DR detach, preparation for, 22

DR-unsafe device, 18

drvconfig command for reconfiguring devices, 40

drvconfig, for reconfiguring devices, 43, 45

N

network device detach, preparation for, 22

non-network devices and the detach command, 54

U

unsafe device, 18

