



Sun Fire Midrange Systems Dynamic Reconfiguration User Guide

Sun Microsystems, Inc.
www.sun.com

Part No. 817-4585-10
February 2004, Revision A

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

Copyright 2003, 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

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

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

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

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

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

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

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

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

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

Copyright 2003, 2004 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, Etats-Unis. Tous droits réservés.

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

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

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

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

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

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

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

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

Contents

Preface	vii
Before You Read This Book	vii
How This Book Is Organized	vii
Using UNIX Commands	viii
Shell Prompts	viii
Typographic Conventions	ix
Related Documentation	ix
Accessing Sun Documentation	x
Contacting Sun Technical Support	x
Sun Welcomes Your Comments	x
1. Introduction to DR on Sun Fire Midrange Systems	1
Dynamic Reconfiguration	1
Command Line Interface	2
Graphical User Interface	2
DR Concepts	3
Detachability	3
Quiescence	3
Suspend-Safe and Suspend-Unsafe Devices	4
Attachment Points	4

DR Operations	5
Hot-Plug Hardware	6
Conditions and States	6
Board States and Conditions	6
Board Receptacle States	6
Board Occupant States	7
Board Conditions	7
Component States and Conditions	8
Component Receptacle States	8
Component Occupant States	8
Component Conditions	8
Component Types	9
Sun Fire Midrange System Domains	9
DR on I/O Boards	10
Nonpermanent and Permanent Memory	11
Target Memory Constraints	11
An Illustration of DR Concepts	11
Limitations	13
Memory Interleaving	13
Reconfiguring Permanent Memory	14
2. Command Line Interface	15
The <code>cfgadm</code> Command	16
Displaying Basic Board Status	16
Displaying Detailed Board Status	17
Command Options	19
Testing Boards and Assemblies	20
Installing or Replacing Boards	23
Hot-Swapping a CompactPCI Card	28

3. Troubleshooting	33
Unconfigure Operation Failure	33
CPU/Memory Board Unconfiguration Failures	34
Cannot Unconfigure a Board Whose Memory Is Interleaved Across Boards	34
Cannot Unconfigure a CPU to Which a Process is Bound	34
Cannot Unconfigure a CPU Before All Memory is Unconfigured	35
Unable to Unconfigure Memory on a Board With Permanent Memory	35
Unable to Unconfigure a CPU	36
Unable to Disconnect a Board	37
I/O Board Unconfiguration Failure	37
Device Busy	37
Problems with I/O Devices	38
RPC or TCP Time-out or Loss of Connection	39
Configure Operation Failure	39
CPU/Memory Board Configuration Failure	39
Cannot Configure Either CPU0 or CPU1 While the Other Is Configured	39
CPUs on a Board Must Be Configured Before Memory	39
I/O Board Configuration Failure	40
Glossary	41

Preface

This book describes the Dynamic Reconfiguration (DR) feature of the Sun™ Fire midrange systems, which include Sun Fire server models 6800, 4810, 4800 and 3800. DR enables you to attach and detach system boards from a running system.

Before You Read This Book

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

How This Book Is Organized

This book contains the following chapters:

Chapter 1 Introduction to DR on Sun Fire Midrange Systems

Chapter 2 Command Line Interface

Chapter 3 Troubleshooting

Glossary

Using UNIX Commands

This document might not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices. See the following for this information:

- Software documentation that you received with your system
- Solaris™ operating environment documentation, which is at <http://docs.sun.com>

Shell Prompts

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

Typographic Conventions

Typeface ¹	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. Replace command-line variables with real names or values.	Read Chapter 6 in the <i>User's Guide</i> . These are called <i>class</i> options. You <i>must</i> be super user to do this. To delete a file, type <code>rm filename</code> .

¹ The settings on your browser might differ from these settings.

Related Documentation

TABLE P-1 Related Documentation

Application	Title
Platform administration	<i>Sun Fire Midrange Systems Platform Administration Manual</i>
System Controller commands	<i>Sun Fire Midrange System Controller Command Reference Manual</i>
Firmware release notes	<i>Sun Fire Midrange Systems Firmware Release Notes</i>
Service Manual	<i>Sun Fire Midrange Systems Service Manual</i>
Internet Multipathing (IPMP)	<i>IP Network Multipathing Administration Guide</i>
Sun Management Center software	<i>Sun Management Center Software User Guide</i>

Accessing Sun Documentation

You can view, print, or purchase a broad selection of Sun documentation, including localized versions, at:

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

Contacting Sun Technical Support

If you have technical questions about this product that are not answered in this document, go to:

<http://www.sun.com/service/contacting>

Sun Welcomes Your Comments

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

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

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

Sun Fire Midrange Systems Dynamic Reconfiguration User Guide, part number 817-4585-10

Introduction to DR on Sun Fire Midrange Systems

While the Solaris 9 operating system supports the full functionality of DR, some previous versions of the Solaris operating environment did not support reconfiguration of I/O boards.

Solaris 8 2/02 software is the first release of the Solaris 8 operating system to support the full functionality of DR on domains. Requirements include appropriate patches and a new kernel update on the domain.

For complete information and instructions for enabling DR on a system running Solaris 8 software, visit:

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

Note – Performing DR operations requires root access.

Dynamic Reconfiguration

DR software is part of the Solaris operating environment. With DR software you can dynamically reconfigure system boards and safely remove them or install them into a system while the Solaris operating environment is running and with minimum disruption to user processes running in the domain.

You can use DR to do the following:

- Minimize the interruption of system applications while installing or removing a board.
- Disable a failing device by removing it from the domain, before the failure can crash the operating system.

- Display the operational status of boards in a domain.
- Initiate system tests of a board while the system continues to run.
- Reconfigure a domain while Solaris continues to run in the domain.
- Invoke hardware-specific functions of a board or a related attachment.

Command Line Interface

The DR software has a command line interface (CLI) using the `cfgadm` command, which is the configuration administration program. The DR agent also provides a remote interface to Sun™ Management Center software.

Graphical User Interface

The optional Sun Management Center software (beginning with version 3.0) provides features such as domain management, as well as a graphical user interface (GUI) to the `cfgadm` DR command line interface (CLI). If you prefer to use a GUI, use Sun Management Center software instead of the command line interfaces of the system controller software and the DR software.

To use Sun Management Center software, you must attach the System Controller board to a network. With a network connection, you can view both the command line interface and the graphical user interface. For instructions on how to use Sun Management Center software, refer to the *Sun Management Center User's Guide*, shipped with Sun Management Center software. For instructions on how to connect the system controller to a network connection on the System Controller board, refer to your systems installation documentation.

DR Concepts

This section contains descriptions of general DR concepts that pertain to Sun Fire midrange system domains.

Detachability

For a device to be detachable, it must conform to the following items:

- Device driver must support `DDI_DETACH`.
- Critical resources must be redundant or accessible through multiple pathways. CPUs and memory banks can be redundant critical resources. Disk drives are examples of critical resources that can be accessible through multiple pathways.

Some boards cannot be detached because their resources cannot be moved. For example, if a domain has only one CPU board, that CPU board cannot be detached. If the boot drive does not have the failover feature implemented, the I/O board connected to it is not detachable.

If there are not multiple pathways 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 so that the I/O board can be detached without losing access to the secondary disk chain.

Quiescence

During the unconfigure operation on a system board with permanent memory (OpenBoot™ PROM or kernel memory), the operating environment is briefly paused, which is known as operating environment quiescence. All operating environment and device activity on the centerplane must cease during a critical phase of the operation.

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:

- An execution 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.

Suspend-Safe and Suspend-Unsafe Devices

When DR suspends the operating environment, all of the device drivers that are attached to the operating environment must also be suspended. If a driver cannot be suspended (or subsequently resumed), the DR operation fails.

A *suspend-safe* device 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.

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

Attachment Points

An attachment point is a collective term for a board and its 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*) has 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(1M)` command.
- An occupant I/O board includes any external storage devices connected by interface cables.

There are two formats used when referring to attachment points:

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

```

/devices/ssm@0,0:N0.SBx    (for a CPU/Memory board)
OR
/devices/ssm@0,0:N0.IBx    (for an I/O assembly)

```

where N0 is node 0 (zero),

SB is a system board,

IB is an I/O board, and

x is a slot number. A slot number can range from 0 through 5 for a system board, and from 6 through 9 for an I/O board.

- A *logical* attachment point is an abbreviated name created by the system to refer to the physical attachment point. Logical attachment points take one of the following two forms:

```

N0.SBx    (for a CPU/Memory board)
OR
N0.IBx    (for an I/O assembly)

```

DR Operations

There are four main types of DR operations.

Operation	Description
Connect	The slot provides power to the board and monitors its temperature. For I/O boards, the connection operation is included in the configuration operation.
Configure	The operating environment assigns functional roles to a board, and loads device drivers for the board and for devices attached to the board.
Unconfigure	The system detaches a board logically from the operating environment and takes the associated device drivers offline. Environmental monitoring continues, but devices on the board are not available for system use.
Disconnect	The system stops monitoring the board, and power to the slot is turned off.

If a system board is in use, stop its use and disconnect it from the domain before you power it off. 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.

The `cfgadm(1M)` command 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 have hot-plug connectors can be inserted or removed while the system is running.

I/O boards and CPU/Memory boards used in the Sun Fire midrange servers are hot-plug devices. Some devices, such as the peripheral power supply, are not hot-plug modules and cannot be removed while the system is running.

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.

Before you attempt to perform any DR operation on a board or component from a domain, you must determine state and condition. Use the `cfgadm(1M)` command with the `-la` options to display the type, state, and condition of each component and the state and condition of each board slot in the domain. See the section [“Component Types” on page 9](#) for a list of the component types.

Board States and Conditions

This section contains descriptions of the states and conditions of system boards (also known as system slots).

Board Receptacle States

A board can have one of three receptacle states: empty, disconnected, or connected. Whenever you insert a board, the receptacle state changes from empty to disconnected. Whenever you remove a board the receptacle state changes from disconnected to empty.



Caution – Physically removing a board that is in the connected state, or that is powered on and in the disconnected state, crashes the operating system and can result in permanent damage to that system board.

Name	Description
empty	A board is not present.
disconnected	The board is disconnected from the system bus. A board can be in the disconnected state without being powered off. However, a board must be powered off and in the disconnected state before you remove it from the slot.
connected	The board is powered on and connected to the system bus. You can view the components on a board only after it is in the connected state.

Board Occupant States

A board can have one of two occupant states: configured or unconfigured. The occupant state of a disconnected board is always unconfigured.

Name	Description
configured	At least one component on the board is configured.
unconfigured	All of the components on the board are unconfigured.

Board Conditions

A board can be in one of four conditions: unknown, ok, failed, or unusable.

Name	Description
unknown	The board has not been tested.
ok	The board is operational.
failed	The board failed testing.
unusable	The board slot is unusable.

Component States and Conditions

This section contains descriptions of the states and conditions for components.

Component Receptacle States

A component cannot be individually connected or disconnected. Thus, components can have only one state: connected.

Component Occupant States

A component can have one of two occupant states: configured or unconfigured.

Name	Description
configured	Component is available for use by the Solaris Operating Environment.
unconfigured	Component is not available for use by the Solaris Operating Environment.

Component Conditions

A component can have one of three conditions: unknown, ok, failed.

Name	Description
unknown	Component has not been tested.
ok	Component is operational.
failed	Component failed testing.

Component Types

You can use DR to configure or to unconfigure several types of components. .

Name	Description
cpu	Individual CPU
memory	All the memory on the board
pci	Any I/O device, controller, or bus

Sun Fire Midrange System Domains

The Sun Fire midrange servers can be divided into dynamic system domains, referred to as *domains* in this document. These domains are based on system board slots that are assigned to the domains. Each domain is electrically isolated into hardware partitions, which ensures that an arbitrary stop in one domain does not affect the other domains in the server.

The domain configuration is determined by the domain configuration table in the platform configuration database (PCD), which resides on the system controller (SC). The domain table controls how the system board slots are logically partitioned into domains. The domain configuration includes empty slots and populated slots.

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

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

DR on I/O Boards

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

If you need to remove a board with I/O devices from a domain temporarily and then re-add it before any other boards with I/O devices are added or removed, reconfiguration is not necessary and need not be performed. In this case, device paths to the board devices will remain unchanged.

- Run the SC command `showenv` to determine the state and usage of the device.
- If disk mirroring is being used to access a device connected to the board, reconfigure the device so that it is accessible by way of controllers on other system boards.
- Unmount file systems.
- Remove multipathing databases from board-resident partitions. The location of multipathing databases is explicitly chosen by the user and can be changed.

Refer to the version of the *Solaris Release Notes Supplement for Sun Hardware* applicable to your Solaris software release for special instructions for I/O devices.

- Remove any private regions used by volume managers. By default, volume managers use a private region on each device that they control. Such devices must be removed from volume manager control before they can be detached.
- Take any RSM 2000 controllers offline by using the `rm6` or `rdacutil` commands.
- Remove disk partitions from the swap configuration.
- 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.

Note – Either kill any process that directly opens a device or raw partition, or direct it to close the open device on the board. If you use the `ndd(1M)` command to set the configuration parameters for network drivers, the parameters may not persist after a DR operation. Use the `/etc/system` file or the `driver.conf` file for a specific driver to set the parameters permanently.

Nonpermanent and Permanent Memory

Before you can delete a board, the environment must vacate the memory on that board. Vacating a board means flushing its nonpermanent memory to swap space and copying its permanent (that is, kernel and OpenBoot™ PROM memory) to another memory board. To relocate permanent memory, the operating environment on a domain must be temporarily suspended, or quiesced. The length of the suspension depends on the domain I/O configuration and the running workloads. Detaching a board with permanent memory is the only time when the operating environment is suspended; therefore, you should know where permanent memory resides so that you can avoid significantly impacting the operation of the domain. You can display the permanent memory by using the `cfgadm(1M)` command with the `-v` option. When permanent memory is on the board, the operating environment must find another memory component of adequate size to receive the permanent memory.

Target Memory Constraints

When permanent memory is removed, DR chooses a target memory area to receive a copy of the memory. The DR software automatically checks for total adherence. It does not allow the DR memory operation to continue if it cannot verify total adherence. A DR memory operation can be disallowed because the domain does not have enough available memory to hold the permanent memory.

An Illustration of DR Concepts

DR lets you disconnect and then reconnect system boards without bringing the system down. You can use DR to add or remove system resources while the system continues to operate.

As an example reconfiguration of system resources, consider the following Sun Fire system configuration, as depicted in the diagram that follows: domain A contains system boards 0 and 2, and I/O board 7. Domain B contains system boards 1 and 3, and I/O board 8.

Note – Before performing DR operations, always make sure that the system complies with the constraints set forth in [“Limitations” on page 13](#).




FIGURE 1-1 Example Domains Before Reconfiguration

To re-assign system board 1 from domain B to domain A, you can use the Sun Management Center software GUI. Or you can perform the following steps manually on the CLI in each domain:

1. As super user, enter the following command on the command line in domain B to disconnect system board 1:

```
# cfgadm -c disconnect -o unassign N0.SB1
```

2. Then, enter the following command on the command line in domain A to assign, connect, and configure system board 1 in Domain A:

```
# cfgadm -c configure N0.SB1
```

The following system configuration is the result. Notice that only the way in which the boards are connected has changed, but not the physical layout of the boards within the cabinet.




FIGURE 1-2 Example Domains After Configuration

Limitations

Memory Interleaving

System boards cannot be dynamically reconfigured if system memory is interleaved across multiple CPU/Memory boards.

Note – For more information about memory interleaving, refer to the `interleave-scope` parameter of the `setupdomain` command, which is described in both the *Sun Fire Midrange Systems Platform Administration Manual* and the *Sun Fire Midrange System Controller Command Reference Manual*.

Conversely, CompactPCI cards and I/O boards *can* be dynamically reconfigured whether memory is interleaved or not.

Reconfiguring Permanent Memory

When a CPU/Memory board containing non-relocatable (permanent) memory is dynamically reconfigured out of the system, a short pause in all domain activity is required which may delay application response. Typically, this condition applies to one CPU/Memory board in the system. The memory on the board is identified by a non-zero permanent memory size in the status display produced by the `cfgadm -av` command.

DR supports reconfiguration of permanent memory from one system board to another only if one of the following conditions is met:

- The target system board has the same amount of memory as the source system board;
- OR-
- The target system board has more memory than the source system board. In this case, the additional memory is added to the pool of available memory.

Command Line Interface

The following procedures are discussed in this chapter:

- [“To Test an I/O Assembly” on page 21](#)
- [“To Install a New Board in a Domain” on page 23](#)
- [“To Hot-Swap a CPU/Memory Board” on page 24](#)
- [“To Hot-Swap an I/O Assembly” on page 25](#)
- [“Hot-Swapping a CompactPCI Card” on page 28](#)
- [“To Hot-Plug a CompactPCI Card” on page 29](#)
- [“To Remove a Board From the System” on page 30](#)
- [“To Move a Board Between Domains” on page 31](#)
- [“To Disconnect a Board Temporarily” on page 32](#)

Note – There is no need to enable dynamic reconfiguration explicitly on Sun Fire midrange systems, as DR is enabled by default. However, for full functionality on systems running Solaris 8, a later kernel update and certain patches are required. For more information, please visit the following Sun websites:

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

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

The `cfgadm` Command

The `cfgadm(1M)` command provides configuration administration operations on dynamically reconfigurable hardware resources. The following table lists the DR board states.

TABLE 2-1 DR Board States from the System Controller (SC)

Board States	Description
Available	The slot is not assigned to any particular domain.
Assigned	The board belongs to a domain, but the hardware has not been configured to use it. The board may be reassigned by the chassis port or released by the domain that it is assigned to.
Active	The board is being actively used by the domain to which it has been assigned. You cannot reassign an active board.

Displaying Basic Board Status

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

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 board slots, SCSI buses, and cPCI slots. The following display shows a typical output.

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

```
# cfgadm
Ap_Id          Type          Receptacle  Occupant    Condition
N0.IB6        PCI_I/O_Boa  connected   configured  ok
N0.IB7        PCI_I/O_Boa  connected   configured  ok
N0.IB8        PCI_I/O_Boa  connected   configured  ok
N0.IB9        PCI_I/O_Boa  disconnected unconfigured unknown
N0.SB0        CPU_Board    connected   configured  unknown
N0.SB1        CPU_Board    disconnected unconfigured failed
```

CODE EXAMPLE 2-1 Output of the Basic `cfgadm` Command (*Continued*)

Displaying Detailed Board Status

For a more detailed status report, use the command `cfgadm -av`. The `-a` option lists attachment points and the `-v` option turns on expanded (verbose) descriptions.

CODE EXAMPLE 2-2 is a *partial* display produced by the `cfgadm -av` command. The output appears complicated because the lines wrap around in this display. (This status report is for the same system shown on [page 19](#), and provides details of each display item.)

CODE EXAMPLE 2-2 Output of the `cfgadm -av` Command

CODE EXAMPLE 2-2 Output of the `cfgadm -av` Command

```

N0.IB7::pci1      connected      configured      ok              device
/ssm@0,0/pci@1b,600000
Apr 3 18:04 io          n              /devices/ssm@0,0:N0.IB7::pci1
N0.IB7::pci2      connected      configured      ok              device
/ssm@0,0/pci@1a,700000
Apr 3 18:04 io          n              /devices/ssm@0,0:N0.IB7::pci2
N0.IB7::pci3      connected      configured      ok              device
/ssm@0,0/pci@1a,600000
Apr 3 18:04 io          n              /devices/ssm@0,0:N0.IB7::pci3
N0.IB8            connected      configured      ok              powered-on, assigned
Apr 3 18:04 PCI_I/O_Boa n              /devices/ssm@0,0:N0.IB8
N0.IB8::pci0      connected      configured      ok              device
/ssm@0,0/pci@1d,700000
Apr 3 18:04 io          n              /devices/ssm@0,0:N0.IB8::pci0
N0.IB8::pci1      connected      configured      ok              device
/ssm@0,0/pci@1d,600000
Apr 3 18:04 io          n              /devices/ssm@0,0:N0.IB8::pci1
N0.IB8::pci2      connected      configured      ok              device
/ssm@0,0/pci@1c,700000, referenced
Apr 3 18:04 io          n              /devices/ssm@0,0:N0.IB8::pci2
N0.IB8::pci3      connected      configured      ok              device
/ssm@0,0/pci@1c,600000, referenced
Apr 3 18:04 io          n              /devices/ssm@0,0:N0.IB8::pci3
N0.IB9            disconnected   unconfigured    unknown         powered-on, assigned
Apr 3 18:04 PCI_I/O_Boa n              /devices/ssm@0,0:N0.IB9
N0.SB0            connected      configured      unknown         powered-on, assigned
Apr 3 18:04 CPU_Board  n              /devices/ssm@0,0:N0.SB0
N0.SB0::cpu0      connected      configured      ok              cpuid 0, speed 750 MHz,
ecache 8 MBytes
Apr 3 18:04 cpu          n              /devices/ssm@0,0:N0.SB0::cpu0
N0.SB0::cpu1      connected      configured      ok              cpuid 1, speed 750 MHz,
ecache 8 MBytes
Apr 3 18:04 cpu          n              /devices/ssm@0,0:N0.SB0::cpu1
N0.SB0::cpu2      connected      configured      ok              cpuid 2, speed 750 MHz,
ecache 8 MBytes
Apr 3 18:04 cpu          n              /devices/ssm@0,0:N0.SB0::cpu2

```

Here are some details of the previous display:

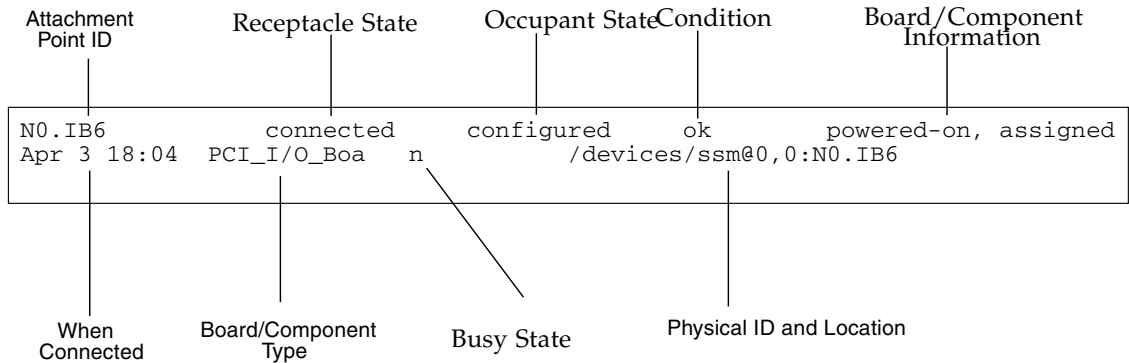


FIGURE 2-1 Details of the Display for `cfgadm -av`

Command Options

The options to the `cfgadm -c` command are listed below:

TABLE 2-2 `cfgadm -c` Command Options

cfgadm -c Option	Function
<code>connect</code>	The slot provides power to the board and begins monitoring the board. The slot is assigned if it was not previously assigned.
<code>disconnect</code>	The system stops monitoring the board and power to the slot is turned off.
<code>configure</code>	The operating system assigns functional roles to a board and loads device drivers for the board and for the devices attached to the board.
<code>unconfigure</code>	The system detaches a board logically from the operating system and takes the associated device drivers offline. Environmental monitoring continues, but any devices on the board are not available for system use.

The options provided by the `cfgadm -x` command are listed below.

TABLE 2-3 `cfgadm -x` Command Options

cfgadm -x Option	Function
<code>assign</code>	Adds (assigns) a board to a domain.
<code>unassign</code>	Deletes (unassigns) a board from a domain.
<code>poweron</code>	Powers on a system board.
<code>poweroff</code>	Powers off a system board.

The `cfgadm_sbd` man page provides additional information on the `cfgadm -c` and `cfgadm -x` options. The `sbd` library provides the functionality for hot-plugging system boards of the class `sbd`, through the `cfgadm` framework.

Testing Boards and Assemblies

▼ To Test a CPU/Memory Board

Before you can test a CPU/Memory board, it must first be assigned to a domain, powered-on, and disconnected. If all these conditions are not met, the board test fails.

You can use the Solaris `cfgadm` command to test CPU/memory boards. As super user type:

```
# cfgadm -t ap-id
```

To change the level of diagnostics that `cfgadm` runs, supply a diagnostic level for the `cfgadm` command as follows.

```
# cfgadm -o platform=diag=<level> -t ap-id
```

where *level* is a diagnostic level, and *ap-id* is an attachment point identifier.

If you do not supply level, the default diagnostic level is set by the `setupdomain` command, which is described in both the *Sun Fire Midrange Systems Platform Administration Manual* and the *Sun Fire Midrange System Controller Command Reference Manual*. The diagnostic levels are:

TABLE 2-4 Diagnostic Levels

Diagnostic Level	Description
init	Only system board initialization code is run. No testing is done. This is a very fast pass through the POST.
quick	All system board components are tested with few tests and test patterns.
default	All system board components are tested with all tests and test patterns, except for memory and Ecache modules. Note that max and default are the same definition.
max	All system board components are tested with all tests and test patterns, except for memory and Ecache modules. Note that max and default are the same definition.
mem1	Runs all tests at the default level, plus more exhaustive DRAM and SRAM test algorithms. For Memory and Ecache modules, all locations are tested with multiple patterns. More extensive, time-consuming algorithms are not run at this level.
mem2	The same as mem1, with the addition of a DRAM test that does explicit compare operations of the DRAM data.

▼ To Test an I/O Assembly

An I/O assembly should be tested before it is added to a domain.

In order to test an I/O assembly you must have a spare domain that is not running the Solaris operating environment.

1. Enter the domain shell of a spare domain (A-D) that is *not* running the Solaris operating environment and that has at least one CPU/Memory board.
2. Press and hold the CTRL key while pressing the] key to bring up the `telnet>` prompt, then type `send break` to display the system controller domain shell.

Note – In this example, domain A is the current, active domain; and domain B is the spare domain.

3. In the spare domain (B) shell, add the I/O assembly to the domain with the `addboard` command.

```
schostname:B> addboard IBx
```

where *x* is 6, 7, 8, or 9.

4. Set the virtual keyswitch in the spare domain to on.

```
schostname:B> setkeyswitch on
.
.
{x} ok
```

where *x* represents the CPU. POST is run on the domain when you turn the virtual keyswitch to on. If you see the `ok` prompt, the I/O assembly is functioning properly.

5. Type:

```
schostname:B> setkeyswitch standby
```

6. Delete the board by entering:

```
schostname:B> deleteboard ibx
```

7. On the active domain (A) add the board using the following command:

```
# cfgadm -c configure N0.IBx
```

Installing or Replacing Boards

▼ To Install a New Board in a Domain



Caution – For complete information about physically removing and replacing boards, refer to the *Sun Fire Midrange Systems Service Manual*. Failure to follow the stated procedures can result in damage to system boards and other components. Also refer to the *Sun Fire Midrange Systems Platform Administration Manual* for more information about software procedures related to removing and replacing boards and components.

Note – When replacing boards, you sometimes need filler panels. A fully configured Sun Fire midrange system ships with three different filler panels: one system board filler panel, one CompactPCI filler panel, and one L2 Repeater Board filler panel.

If you are unfamiliar with how to insert a board into the system, refer to the *Sun Fire Midrange Systems Service Manual* before you begin this procedure.

1. **Identify an empty slot available to the domain by typing the following as super user:**

```
# cfgadm -l -s "select=class(sbd)"
```

2. **Make sure you are properly grounded with a wrist strap.**
3. **After locating the empty slot, remove the system board filler panel from the slot.**
4. **Insert the board into the slot** within one minute to prevent system overheating. Refer to the *Sun Fire Midrange Systems Service Manual* for complete step-by-step board insertion procedures.

5. Power on, test, and configure the board using the `cfgadm -c configure` command.

```
# cfgadm -c configure ap_id
```

where *ap_id* is the attachment point ID returned by `cfgadm -l -s "select=class(sbd)"`.

▼ To Hot-Swap a CPU/Memory Board



Caution – For complete information about physically removing and replacing boards, refer to the *Sun Fire Midrange Systems Service Manual*. Failure to follow the stated procedures can result in damage to system boards and other components.

Note – Hot-swapping is initiated by the user by pressing the card's ejector lever fully while the card is inserted; or by disengaging the ejector lever partially before the card is removed. The operator does not need to issue any commands to perform a hot swap. Hot-plugging, on the other hand, is accomplished using the `cfgadm` command.

1. If the board is being used by the Solaris operating environment, as super user identify the board to be removed.

You must know the slot number (attachment point ID).

```
# cfgadm -l -s "select=class(sbd)"
```



2. Make sure you are properly grounded using a wrist strap.
3. Detach the board from the domain and power off the board with `cfgadm`.

```
# cfgadm -c disconnect ap_id
```

where *ap_id* is the attachment point ID.

This command removes the resources from the Solaris operating environment and the OpenBoot PROM, detaches the board from the domain, and powers off the board.

4. **Verify the state of the Power and Hotplug OK LEDs.**

The green Power LED will flash shortly as the CPU/Memory board is cooling down. In order to safely remove the board from the systems the green Power LED  must be off and the amber Hotplug OK LED  must be on.


5. **Complete the hardware removal and installation of the board. For more information refer to the *Sun Fire Midrange Systems Service Manual*.**

6. **After removing and installing board, bring the board back to the Solaris operating environment with the Solaris dynamic reconfiguration `cfgadm` command.**

```
# cfgadm -c configure ap_id
```

where *ap_id* is the attachment point ID.

This command assigns the board to the domain, powers it on, tests it, attaches the board, and brings all of its resources back to the Solaris operating environment.

7. **Verify that the green Power LED  is lit.**

▼ To Hot-Swap an I/O Assembly

There are two types of I/O assemblies: CompactPCI (cPCI) and standard PCI. These instructions apply to both types. Note, however, that while cPCI cards can be hot-swapped, hot-plugged, and dynamically re-configured, PCI cards and standard I/O assemblies cannot be hot-swapped, hot-plugged, nor dynamically reconfigured.

Hot-swapping is initiated by the user by pressing the card's ejector lever fully while the card is inserted; or by disengaging the ejector lever partially before the card is removed. The operator does not need to issue any commands to perform a hot swap. Hot-plugging, on the other hand, is accomplished using the `cfgadm` command.



Caution – For complete information about physically removing and replacing boards, refer to the *Sun Fire Midrange Systems Service Manual*. Failure to follow the stated procedures can result in damage to system boards and other components.

1. **If the I/O assembly is being used by the Solaris operating environment, as super user in the Solaris operating environment, identify the I/O assembly to be removed.**

You must know the slot number (attachment point ID).

```
# cfgadm -l -s "select=class(sbd)"
```

2. Detach the board from the domain and power off the board with `cfgadm`:

```
# cfgadm -c disconnect ap_id
```



Where: *ap_id* is the attachment point ID.

This command removes the resources from the Solaris operating environment and the OpenBoot PROM, detaches the board from the domain, and powers off the I/O assembly.

3. Remove the board from the domain with `cfgadm`:

```
# cfgadm -x unassign ap_id
```

4. Verify the state of the status LEDs on the I/O assembly.

In order to safely remove the I/O assembly from the system, the green Power LED  on the I/O assembly must be in the deactivated state (off) and the amber Hotplug OK LED  must be lit.

5. Complete the hardware removal and installation of the I/O assembly. For more information see the *Sun Fire Midrange Systems Service Manual*.

Note – Be sure you are properly grounded before you begin the hardware removal and installation of the I/O assembly.

Before you bring the board back to the Solaris operating environment, you need to enter a spare domain that is NOT running the Solaris operating environment and that has at least one CPU/Memory board in order test the I/O assembly.

Enter the domain shell of a spare domain (A-D) that is NOT running the Solaris operating environment and that has *at least* one CPU/Memory board.

6. Press and hold the CTRL key while pressing the] key to bring up the `telnet>` prompt. Type `send break` to display the system controller domain shell.

Note – In this example, domain A is the current, active domain; and domain B is used as a spare domain.

7. In the spare domain shell, add the I/O assembly to the domain with the `addboard` command.

```
schostname:B> addboard ibx
```

where x is 6, 7, 8, or 9.

8. Set the virtual keyswitch in the spare domain to on.

POST is run on the domain when you turn the virtual keyswitch to on.

```
schostname:B> setkeyswitch on
.
.
{x} ok
```

where x represents the CPU. If you see the `ok` prompt, the I/O assembly is functioning properly.

9. Press and hold the CTRL key while pressing the J key to bring up the `telnet>` prompt. Type `send break` to connect to the system controller domain shell.

Depending on the type of telnet connection, you may need to type `send esc` followed by `send break` to connect to the system controller domain shell.

10. Type:

```
schostname:B> setk standby
```

11. Delete the board by entering:

```
schostname:B> deleteboard ibx
```

12. At the Solaris prompt in domain A configure the I/O assembly:

```
#cfgadm -c configure N0.IBx
```

Hot-Swapping a CompactPCI Card

You can initiate hot-swapping by pressing the card's ejector lever fully while the card is inserted; or by disengaging the ejector lever partially before the card is removed. You do not need to issue any commands to perform a hot swap. To perform a hot-plug operation, on the other hand, use the `cfgadm` command.

In order to hot-swap a CompactPCI (cPCI) card, you must boot the Solaris software in the domain where the cPCI card I/O assembly resides. When the Solaris software is booted in the domain, all cPCI cards are in the autoconfigure mode, and all configuring and unconfiguring can be performed without the `cfgadm` command.

When you insert a cPCI card using hot-swap, the card is automatically powered on and configured. When you remove a cPCI card using hot-swap, the card is automatically unconfigured and powered off.




Caution – For complete information about physically removing and replacing boards, refer to the *Sun Fire Midrange Systems Service Manual*. Failure to follow the stated procedures can result in damage to system boards and other components.

▼ To Insert a CompactPCI Card

1. As super user, identify the slot into which the card will be inserted.
2. Insert the card and push down on the ejector lever fully to engage it reliably. The card will be automatically powered-on and configured. The blue Hotswap OK LED on the card should be off, the green Power LED on the I/O assembly should be lit, and the amber Hotplug OK LED should be off. Insertion using hot-swap is equivalent to typing the following command: `cfgadm -c configure ap_id`.

▼ To Remove a CompactPCI Card

Note – Before you hot-swap the CompactPCI (cPCI) card, make sure that there is no I/O activity on that card.

1. Disengage the ejector lever slightly to deactivate the card.
2. Make sure the blue Hotswap OK LED on the card is lit, the amber Hotplug OK LED  on the I/O assembly is lit, and the green Power LED on the assembly is off.

3. Remove the card.

If the domain console is available, a message confirms that the card has been unconfigured.

▼ To Hot-Plug a CompactPCI Card

Hot-plugging is accomplished by using the `cfgadm` command. You perform a hot-swap operation, on the other hand, by pressing the card's ejector lever fully while the card is inserted; or by disengaging the ejector lever partially before the card is removed. You never need to issue any commands to perform a hot swap.

1. As super user, identify the cPCI card to be removed.

You must know the slot number (attachment point ID).

```
# cfgadm -s "select=class(pci)"
```

2. Detach (unconfigure) the cPCI card to be removed:

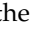
```
# cfgadm -c unconfigure ap_id
```

where `ap_id` is the attachment point ID. The card is automatically unconfigured and powered-off.

3. Confirm that the card is detached.

```
# cfgadm -s "select=class(pci)" ap_id
```

4. Inspect the green Power LED and the amber Hotplug OK LED on the I/O assembly and the blue Hotswap OK LED on the cPCI card.

When the green Power LED on the I/O assembly is off, the amber Hotplug OK LED  on the I/O assembly is lit, and the blue Hotswap OK LED on the cPCI card if lit, it is safe to remove the cPCI card.

5. After ensuring that you are properly grounded using a wrist strap, remove and replace the cPCI card.



Caution – For complete information about physically removing and replacing boards, refer to the *Sun Fire Midrange Systems Service Manual*. Failure to follow the stated procedures can result in damage to system boards and other components.


6. After installing the card, attach (configure) the card:

```
# cfgadm -c configure ap_id
```

When Solaris boots in a domain, all cPCI slots in that domain are in the autoconfigure mode by default. In autoconfigure mode hot-swap is enabled for each slot.

Note – To disable the Autoconfiguration feature, use the following command: **cfgadm -x disable_autoconfig ap_id**. To re-enable Autoconfiguration, use the following command: **cfgadm -x enable_autoconfig ap_id**.

7. Inspect the green Power LED.

The green Power LED on the I/O assembly  will be lit and the blue Hotswap OK LED on the cPCI card should be off.

8. Verify that the card is attached.

```
# cfgadm -s "select=class(pci)" ap_id
```

▼ To Remove a Board From the System

Note – Before you begin this procedure, make sure you have ready a system board filler panel to replace the system board you are going to remove from the system. A system board filler panel is a metal board with slots that allow cooling air to circulate.

1. Identify the board to be removed.

You must know the slot number.

```
# cfgadm -l -s "select=class(sbd)"
```


2. Detach and power off the board from the domain by using the `cfgadm -c disconnect` command.

```
# cfgadm -c disconnect ap_id
```

where *ap_id* is the attachment point ID returned by `cfgadm -al -s select=class(sbd)`.



Caution – For complete information about physically removing and replacing boards, refer to the *Sun Fire Midrange Systems Service Manual*. Failure to follow the stated procedures can result in damage to system boards and other components.

3. Remove the board from the system.

Refer to the *Sun Fire Midrange Systems Service Manual* for complete step-by-step board removal procedures.

4. Insert a system board filler panel into the slot within one minute of removing the board to prevent system overheating.

▼ To Move a Board Between Domains

1. Identify the slot number of the board to be removed.

```
# cfgadm -l -s "select=class(sbd)"
```

2. Unconfigure the board but leave the power on to preserve the test status:

```
# cfgadm -o unassign,nopoweroff -c disconnect ap_id
```

where *ap_id* is the attachment point ID returned by `cfgadm -l -s "select=class(sbd)"`.

At this point, the slot is not assigned to any domain, and the slot is visible to all domains.

3. In the domain to which you are moving the board, check to see if the board is now visible as disconnected.

```
# cfgadm -al -s select=class(sbd)
```

4. Configure the board in the new domain using the `cfgadm -c configure` command, which implies an assignment operation.

```
# cfgadm -c configure ap_id
```

▼ To Disconnect a Board Temporarily

You can use DR to power down the board and leave it in place. For example, you might want to do this if the board fails and a replacement board or a system board filler panel is not available.

1. Identify the board to be removed.

You must know the slot number.

```
# cfgadm -l -s "select=class(sbd)"
```

2. Detach and power off the board from the domain by using the `cfgadm -c disconnect` command.

```
# cfgadm -c disconnect ap_id
```

where *ap_id* is the attachment point ID returned by `cfgadm -l -s "select=class(sbd)"`.

Troubleshooting

This chapter discusses common types of failure:

- [“Unconfigure Operation Failure” on page 33](#)
- [“Configure Operation Failure” on page 39](#)

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

See the following man pages for additional error message detail: `cfgadm(1M)`, `cfgadm_sbd(1M)`, `cfgadm_pci(1M)`, and `config_admin(3X)`.

Unconfigure Operation Failure

An unconfigure operation for a CPU/Memory board or an I/O board can fail if the system is not in a correct state before you begin the operation.

CPU/Memory Board Unconfiguration Failures

- Memory on a board is interleaved across boards before an attempt to unconfigure the board.
- A process is bound to a CPU before an attempt to unconfigure the CPU.
- Memory remains configured on a system board before you attempt a CPU unconfigure operation on that board.
- The memory on the board is configured (in use). See [“Unable to Unconfigure Memory on a Board With Permanent Memory”](#) on page 35.
- CPUs on the board cannot be taken off line. See [“Unable to Unconfigure a CPU”](#) on page 36.

Cannot Unconfigure a Board Whose Memory Is Interleaved Across Boards

If you try to unconfigure a system board whose memory is interleaved across system boards, the system displays an error message such as:

```
cfgadm: Hardware specific failure: unconfigure N0.SB2::memory: Memory is
interleaved across boards: /ssm@0,0/memory-controller@b,400000
```

Cannot Unconfigure a CPU to Which a Process is Bound

If you try to unconfigure a CPU to which a process is bound, the system displays an error message such as the following:

```
cfgadm: Hardware specific failure: unconfigure N0.SB2::cpu3: Failed to off-line:
/ssm@0,0/SUNW,UltraSPARC-III
```

- **Unbind the process from the CPU and retry the unconfigure operation.**

Cannot Unconfigure a CPU Before All Memory is Unconfigured

All memory on a system board must be unconfigured before you try to unconfigure a CPU. If you try to unconfigure a CPU before all memory on the board is unconfigured, the system displays an error message such as:

```
cfgadm: Hardware specific failure: unconfigure N0.SB2::cpu0: Can't unconfig cpu
if mem online: /ssm@0,0/memory-controller
```

- **Unconfigure all memory on the board and then unconfigure the CPU.**

Unable to Unconfigure Memory on a Board With Permanent Memory

To unconfigure the memory on a board that has permanent memory, move the permanent memory pages to another board that has enough available memory to hold them. Such an additional board must be available before the unconfigure operation begins.

Memory Cannot Be Reconfigured

If the unconfigure operation fails with a message such as the following, the memory on the board could not be unconfigured:

```
cfgadm: Hardware specific failure: unconfigure N0.SB0: No available memory
target: /ssm@0,0/memory-controller@3,400000
```

Add to another board enough memory to hold the permanent memory pages, and then retry the unconfigure operation.

- **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 -av -s "select=type(memory)"
```

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: unconfigure N0.SB0: Insufficient memory
```

```
cfgadm: Hardware specific failure: unconfigure N0.SB0: Memory operation failed
```

- **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: unconfigure N0.SB0: Memory operation refused
```

- **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 CPU 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 online 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 has not been stopped, or because an I/O device becomes active again after it has been stopped.

Device Busy

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

If a device on the board is busy, the system logs a message such as the following after an attempt to unconfigure:

```
cfgadm: Hardware specific failure: unconfigure N0.IB6: Device
busy: /ssm@0,0/pci@18,700000/pci@1/SUNW,isptwo@4/sd@6,0
```

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

1. To see which processes have these devices open, use the `fuser(1M)` command.
2. Run the following command to kill the vold daemon gracefully.:

```
# /etc/init.d/volmgt stop
```

3. Disconnect all SCSI controllers that are associated with the card that you're trying to unconfigure. To get a list of all connected SCSI controllers use the following command:

```
# cfgadm -l -s "select=class(scsi)"
```

4. If the redundancy features of Solaris™ Volume Manager (SVM) 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.
5. Unmount file systems, including SVM meta-devices that have a board resident partition. (For example, `umount /partition`).
6. Remove the SVM database from board-resident partitions. The location of the SVM database is explicitly chosen by the user and can be changed.
7. 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 Sun Volume Manager control before they can be detached.

8. Remove disk partitions from the swap configuration.
9. Either kill any process that directly opens a device or raw partition, or direct it to close the open device on the board.

Note – 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 Failure

CPU/Memory Board Configuration Failure

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

- You try to configure either CPU0 or CPU1 while the other is configured.
- A CPU remains configured on the board.

Cannot Configure Either CPU0 or CPU1 While the Other Is Configured

Before you try to configure either CPU0 or CPU1, make sure that the other CPU is unconfigured.

CPUs on a Board Must Be Configured Before Memory

Before configuring memory, all CPUs on the system board must be configured. If you try to configure memory while one or more CPUs are unconfigured, the system displays an error message such as:

```
cfgadm: Hardware specific failure: configure N0.SB2::memory: Can't
config memory if not all cpus are online: /ssm@0,0/memory-
controller
```

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. The system logs a message such as the following:

```
cfgadm: Hardware specific failure: configure N0.IB6: Unsafe driver  
present: <device path>
```

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

Glossary

ap_id	Attachment point identifier; an <code>ap_id</code> 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 path name, while a logical identifier contains a shorthand notation.
Attachment point	A collective term for a board and its card cage slot. A <i>physical</i> attachment point describes the software driver and location of the card cage slot. A <i>logical</i> attachment point is an abbreviated name created by the system to refer to the physical attachment point.
cfgadm command	<code>cfgadm</code> is the primary command for dynamic reconfiguration on the Sun Fire midrange systems. For information about the command and its options, refer to the <code>cfgadm(1M)</code> , <code>cfgadm_sbd(1M)</code> , and <code>cfgadm_pci(1M)</code> man pages.
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 <code>DDI_DETACH</code> 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

Domain	A logical grouping of system boards that are electrically connected. Domains are separated from each other and do not interact with one another. Each domain runs its own copy of the Solaris operating environment and has its own host identifier.
Domain administration	The responsibility for connecting and configuring system boards to create domains; and for unconfiguring and disconnecting system boards, either to move them to different domains or to replace defective system boards.
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 IPMP 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.
IP Multipathing (IPMP)	Internet Protocol multipathing. Enables continuous application availability by load balancing failures when multiple network interface cards are attached to a system. If a failure occurs in a network adapter, and if an alternate adapter is connected to the same IP link, the system switches all the network accesses from the failed adapter to the alternate adapter. When multiple network adapters are connected to the same IP link, any increases in network traffic are spread across multiple network adapters, which improves network throughput.
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.
Platform	A specific Sun Fire system model, such as the Sun Fire 6800 system,
Platform administration	The process of setting up domains on a Sun Fire system; re-allocating resources between domains; and monitoring performance on each domain.

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 DDI_SUSPEND 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.
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

A

- active board state, 16
- ADR on I/O boards, 10
- assemblies
 - testing, 20
- assigned board state, 16
- attachment points
 - description of, 4
- Autoconfiguration
 - disabling, 30
 - re-enabling, 30
- available board state, 16
- available component list, 9

B

- board
 - hot-swapping a CPU/memory board, 24
 - installing in a domain, 23
- board state
 - active, 16
 - assigned, 16
 - available, 16
- board status
 - displaying, 16
 - displaying detailed, 17
- boards
 - conditions, 6, 7
 - hot plugging, 6
 - installing or replacing, 23
 - moving between domains, 31

- occupant states, 7
- receptacle states, 6
- removing, 30
- states, 6
- testing, 20
- unconfiguring temporarily, 32

C

- cfgadm
 - cfgadm command, 16
 - cfgadm -v, 17
- cfgadm -c command options, 19
- cfgadm -x command options, 20
- cfgadm(1M)
 - attachment points, 4
 - functions, 6
- CompactPCI card
 - hot-plugging, 29
 - hot-swapping, 28
- component
 - conditions, 8
 - states, 8
 - types, 9
- configured state, 7, 8
- connected state, 7
- CPUs
 - detachability, 3
 - suspending, 3
 - types, 9

D

- DDI_DETACH, 3
- detachability, 3
- disconnected state, 7
- disk
 - mirroring, 10
 - partitions, 10
- domains
 - description of, 9
 - logical, 9
 - physical, 9
 - platform configuration database, 9
- DR
 - concepts, 3
 - operations, 5
- DR-unsafe device, 4
- dynamic reconfiguration (DR)
 - command line interface, 2
 - GUI, 2
 - illustration of concepts, 9, 11
 - introduction, 1
 - limitations, 13
- dynamic system domains, 9

E

- empty
 - slots, 9
 - state, 7

F

- failed condition, 7
- filler panels, 23

H

- hot-plugging boards, 6, 29
- hot-swapping boards, 24, 25, 28

I

- I/O assembly

- hot-swapping, 25

- I/O devices
 - detachability, 3
 - suspending, 3
 - suspend-safe, 4
 - types, 9
 - with ADR, 10

L

- logical attachment point, 5
- logical domain, 9

M

- memory
 - permanent, 11
 - target constraints, 11
- memory types, 9
- multipathing databases, 10

N

- ndd(1M), 10
- Nonpermanent memory, 11

O

- occupant, 4
- ok condition, 7

P

- permanent memory, 11
- physical attachment point, 5
- physical domain, 9
- platform configuration database, 9
- populated slots, 9

Q

- Quiescence, 3

R

- raw partitions, 10
- receptacle, 4
- receptacle state, 6
- RSM 2000 controllers, 10

S

- showdevices(1M)
 - with I/O devices, 10
- slot numbers, 5
- slots, 9
- suspend-safe devices, 4

T

- troubleshooting
 - configure operation failure, 39
 - unconfigure operation, 33

U

- unconfigured state, 7, 8
- unknown condition, 7
- unsafe devices, 4
- unusable condition, 7

V

- volume managers, 10

