



Sun™ Enterprise 10000 DR Configuration Guide

Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95054 U.S.A.
650-960-1300

Part No. 816-3630-10
May 2002, Revision A

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

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 U.S.A. All rights reserved.

This product or document is distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

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

Sun, Sun Microsystems, the Sun logo, AnswerBook2, docs.sun.com, Solstice DiskSuite, Sun FastEthernet, Sun Netra, Sun Enterprise, Sun Enterprise Volume Manager, Sun StorEdge, Sun Ultra, OpenBoot, Solaris, and UltraSparc are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

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

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

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

Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 Etats-Unis. Tous droits réservés.

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

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

Sun, Sun Microsystems, le logo Sun, AnswerBook2, docs.sun.com, Solstice DiskSuite, Sun FastEthernet, Sun Netra, Sun Enterprise, Sun Enterprise Volume Manager, Sun StorEdge, Sun Ultra, OpenBoot, Solaris, et UltraSparc sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

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

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



Adobe PostScript

Sun Enterprise 10000 SSP Attributions:

This software is copyrighted by the Regents of the University of California, Sun Microsystems, Inc., and other parties. The following terms apply to all files associated with the software unless explicitly disclaimed in individual files.

The authors hereby grant permission to use, copy, modify, distribute, and license this software and its documentation for any purpose, provided that existing copyright notices are retained in all copies and that this notice is included verbatim in any distributions. No written agreement, license, or royalty fee is required for any of the authorized uses. Modifications to this software may be copyrighted by their authors and need not follow the licensing terms described here, provided that the new terms are clearly indicated on the first page of each file where they apply.

IN NO EVENT SHALL THE AUTHORS OR DISTRIBUTORS BE LIABLE TO ANY PARTY FOR DIRECT, INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OF THIS SOFTWARE, ITS DOCUMENTATION, OR ANY DERIVATIVES THEREOF, EVEN IF THE AUTHORS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This is scotty, a simple tcl interpreter with some special commands to get information about TCP/IP networks. Copyright (c) 1993, 1994, 1995, J. Schoenwaelder, TU Braunschweig, Germany, Institute for Operating Systems and Computer Networks. Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that this copyright notice appears in all copies. The University of Braunschweig makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

Contents

Preface	vi
Before You Read This Book	vi
How This Book Is Organized	vi
Using UNIX Commands	vii
Typographic Conventions	vii
Shell Prompts	vii
Related Documentation	viii
Accessing Sun Documentation Online	viii
Sun Welcomes Your Comments	ix
1. Configuring DR	1
DR Models	2
Enhancements in DR Model 3.0	3
Where to Execute DR Commands	3
Requirements for Multipathing in DR 3.0	3
Getting Started	3
Device Prerequisites	4
Allocating Sufficient Domain Swap Space	5
Qualifying Third-Party Device Drivers	5

Overview of DR Configuration Tasks 6

- ▼ To Enable the Kernel Cage 7
- ▼ To Set Permanent Driver Parameters for Network Drivers 7
- ▼ To Enable Device Suspension for the `soc` and `pln` Drivers 8
- ▼ To Specify an Unsafe Driver List 8
- ▼ To Make an Unsupported Tape Device Detach-Safe 9

Preparing for DR Detach Operations 9

Configuration Changes During DR Operations 10

Controlling Forcible Conditions that Affect System Quiescence 10

- ▼ To Manually Suspend a Suspend-Unsafe Device 11
- ▼ To Force a System Quiescence 12

Target Memory Constraints 12

Processors 13

Network Devices 13

Remote DR Communication 14

- ▼ To Troubleshoot a Connection Failure During a Solaris 9 (DR Model 3.0) Operation 14

Preface

This guide describes the domain-side configuration of the Sun Enterprise 10000 server Dynamic Reconfiguration (DR) feature. For information about how to use these features, refer to the appropriate document listed in “Related Documentation” on page viii.

Before You Read This Book

This guide is intended for the Sun Enterprise 10000 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 documentation provided with this system, and consider UNIX system administration training.

Also read and be familiar with the *TCP/IP and Data Communications Administration Guide* that is provided with your server.

How This Book Is Organized

This guide contains one chapter that describes how to configure and reconfigure a Sun Enterprise 10000 domain before and after a DR operation.

Using UNIX Commands

This document does not contain complete information on basic UNIX commands and procedures such as shutting down the system, booting the system, and configuring devices. See the Solaris software documentation that you received with your system for this information.

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. 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. To delete a file, type <code>rm filename</code> .

Shell Prompts

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

Related Documentation

Application	Title	Part Number
User	<i>Sun Enterprise 10000 Dynamic Reconfiguration User Guide</i>	816-3627
	<i>Sun Enterprise 10000 SSP 3.5 User Guide</i>	806-7613
	<i>IP Network Multipathing Administration Guide</i>	806-0850
	<i>MPxIO Installation and Configuration Guide</i>	816-1420
	<i>Sun Enterprise 10000 InterDomain Networks User Guide</i>	806-4131
Reference	<i>Sun Enterprise 10000 Dynamic Reconfiguration Reference Manual</i>	806-7617
	<i>Sun Enterprise 10000 SSP 3.5 Reference Manual</i>	806-7614
	<i>Sun Enterprise 10000 IDN Configuration Guide</i>	806-5230
	<i>Sun Enterprise 10000 IDN Error Messages</i>	806-5231
	<i>TCP/IP and Data Communications Administration Guide</i>	806-4003
Installation and Release Notes	<i>Sun Enterprise 10000 SSP 3.5 Installation Guide and Release Notes</i>	806-7615

Accessing Sun Documentation Online

A broad selection of Sun system documentation is located at:

<http://www.sun.com/products-n-solutions/hardware/docs>

A complete set of Solaris documentation and many other titles are located at:

<http://docs.sun.com>

Sun Welcomes Your Comments

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

`docfeedback@sun.com`

Please include the part number (816-3630-10) of your document in the subject line of your email.

Configuring DR

This chapter describes key DR functionality and also guides you through the tasks for configuring DR. You will find

- Information on DR models
- Details on getting started with DR configuration
- An overview of DR configuration tasks
- Prerequisite tasks to be completed before a DR detach operation
- Details on the configuration changes that occur during DR detach operations and how to control certain conditions when a detach operation is in progress

Note – In this document, the phrase “DR detach operation” refers to the complete detach or removal of a system board. This detach operation can be accomplished by using the ADR command `deleteboard(1M)`. For instructions on detaching boards from Solaris 9 domains (which support only DR model 3.0), refer to the *Sun Enterprise 10000 Dynamic Reconfiguration User Guide* (part number 816-3627-10).

DR Models

There are two models of DR available for the Sun Enterprise 10000 system. DR model 2.0 is sometimes referred to as “legacy DR,” and DR model 3.0 is referred to as “next generation DR.” The following table shows the different versions of the Solaris operating environment and the SSP software that are used with DR models 2.0 and 3.0:

DR Model	Solaris Software Versions	SSP Software Versions
2.0	Solaris 5.1, 6, 7, and 8	3.3, 3.4, or 3.5
3.0	Solaris 8 10/01 and 02/02, Solaris 9	3.5 only

Domains running version 9 of the Solaris software support *only* DR model 3.0, and version 3.5 of the SSP software is required.

Only one model of DR can run within a domain at a time. To check the version of DR that is running, use the `domain_status` command with its `-m` option (available only on domains running version 3.5 of the SSP software). Make sure to verify the DR model before you execute any DR commands. The following is an example of the `domain_status(1M)` output. The `DR-MODEL` column indicates which model is enabled

```
# domain_status -m

DOMAIN      TYPE                PLATFORM  DR-MODEL  OS    SYSBDS
A           Ultra-Enterprise-10000  all-A    2.0      5.8   2
B           Ultra-Enterprise-10000  all-A    3.0      5.8   3 4
C           Ultra-Enterprise-10000  all-A    2.0      5.7   5 6
D           Ultra-Enterprise-10000  all-A    3.0      5.9   7
```

According to this output, domain A is running Solaris version 8 software (OS 5.8) with DR model 2.0 enabled; domain B is running Solaris version 8 software with DR model 3.0 enabled; domain C is running Solaris version 7 software (OS 5.7) with DR model 2.0 enabled; and domain D is running Solaris version 9 software (OS 5.9) with DR model 3.0 enabled.

Only certain commands are available in each model, and if you execute a command that is not supported, an error message appears on the console.



Caution – Before you switch to DR 3.0 in a domain that is running the Solaris 8 10/01 operating environment, you must upgrade the SSP software to version 3.5 because previous versions of SSP do not support DR 3.0 operations.

For more information about using DR 2.0, see the the *Sun Enterprise 10000 Dynamic Reconfiguration (DR) User Guide* (part number 806-7616-10). For more information about using DR 3.0, see the *Sun Enterprise 10000 Dynamic Reconfiguration (DR) User Guide* (part number 816-3627-10).

Enhancements in DR Model 3.0

The DR 3.0 model offers the following enhancements to DR 2.0:

- DR 3.0 has a framework that offers better integration with applications, through the Reconfiguration Coordination Manager.
- DR 3.0 supports network multipathing using IPMP.

Where to Execute DR Commands

You execute DR operations from either of two places: from the system service processor (SSP) by using the SSP commands—`addboard(1M)`, `moveboard(1M)`, `deleteboard(1M)`, `rcfgadm(1M)`, and `showdevices(1M)`; or from the domain, using the `cfgadm(1M)` command.

Requirements for Multipathing in DR 3.0

To use multipathing on DR model 3.0 domains, run IPMP (the IP multi-pathing software provided with the Solaris 8 operating environment) and MPxIO software, included in Solaris Kernel Update Patches 111412-02, 111413-02, 111095-02, 111096-02, and 111097-02.

Getting Started

Before you run DR operations on your domain, you must

- Be familiar with how devices must be configured before DR detach operations, as explained in “Device Prerequisites” on page 14.

- Verify that you have sufficient swap space for your domain.
For details, see “Allocating Sufficient Domain Swap Space” on page 15.
- Qualify any third-party device drivers, as described in “Qualifying Third-Party Device Drivers” on page 15.

Device Prerequisites

DR requires that drivers for devices on boards involved in DR detach operations be both:

- *Detach-safe* or not currently loaded

A detach-safe driver supports the device driver interface (DDI) function, `DDI_DETACH`. This function provides the ability to detach a particular instance of a driver without affecting other instances that are servicing other devices.

A *detach-unsafe* driver is one that does not support `DDI_DETACH`. If a detach-unsafe driver is loaded, you must unload it before performing a DR detach operation. For details on unloading a detach-unsafe device, see “Preparing for DR Detach Operations” on page 19.

- *Suspend-safe* or closed

A suspend-safe device driver supports the *quiescence* (pausing) of the Solaris operating environment during the detach of a board that contains nonpageable OBP or kernel memory. In order for DR to perform the detach, the operating environment must temporarily suspend all processes, processors, and device activities to unconfigure the memory component.

A suspend-safe device supports the `DDI_SUSPEND/DDI_RESUME` function. This function enables a device to be suspended during a system quiescence and then resumed. The device managed by the driver will not attempt to access the domain centerplane (for example, it does not access memory or interrupt the system), even if the device is open when the suspend request is made. The quiescence only affects the target domain; other domains are not affected.

If a driver does not support the function `DDI_SUSPEND/DDI_RESUME`, the device is considered to be *suspend-unsafe* because the operating environment cannot quiesce if a suspend-unsafe device is present. If a system quiescence is required for a DR detach operation, you must manually suspend a suspend-unsafe device so that the quiescence can occur. For details, see “To Manually Suspend a Suspend-Unsafe Device” on page 111.

Note – The drivers currently released by Sun Microsystems that are known to be suspend-safe are: `st`, `sd`, `isp`, `esp`, `fas`, `sbus`, `pci`, `pci-pci`, `qfe`, and `hme` (Sun FastEthernet™); `nf` (NPI-FDDI); `qe` (Quad Ethernet); `le` (Lance Ethernet); the SSA drivers (`soc`, `pln`, and `ssd`); and the Sun StorEdge A5000 drivers (`sf`, `socal`, and `ses`). For additional information about suspend-safe and detach-safe device drivers, contact your Sun service representative.

Allocating Sufficient Domain Swap Space

The domain swap configuration consists of the swap devices and `swapfs` (memory). The domain must contain enough swap space so that it can flush pageable memory. For example, if you want to remove 1 Gbyte of memory from a 2-Gbyte domain, you will need 1 Gbyte of swap space, depending on the load. Insufficient swap space can prevent the completion of a DR operation.

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

Note – When memory (`swapfs`) or swap space on a disk is detached, there must be enough memory or swap space remaining in the domain to accommodate currently running programs.

Qualifying Third-Party Device Drivers

Many third-party drivers (those purchased from vendors other than Sun Microsystems) do not support the standard Solaris `modunload(1M)` interface, which is used to unload detach-unsafe or suspend-unsafe device drivers. Conditions that invoke the driver functions occur infrequently during normal operation and the functions are sometimes missing or work improperly. Sun Microsystems suggests that you test these driver functions during the qualification and installation phases of any third-party device.

Overview of DR Configuration Tasks

This section identifies the various configuration tasks that you must complete before running DR operations on Solaris 9 domains (which support only DR model 3.0). Note that it may not be necessary to perform all the tasks described in this section, depending on the types of devices on your system boards and the type of DR operation to be performed.

After you configure DR or whenever you make changes to the DR configuration, you must reboot your domain. If you want to minimize the number of domain reboots, determine which configuration tasks are applicable to your DR environment and then perform the appropriate set of configuration tasks before rebooting your domain.

1. If you intend to perform DR detach operations, enable the kernel cage, as explained in “To Enable the Kernel Cage” on page 17.
2. For devices, do the following:
 - If you set network configuration parameters manually, make these settings permanent as described in “To Set Permanent Driver Parameters for Network Drivers” on page 17.
 - If you have `soc` and `pln` devices, enable device suspension, as described in “To Enable Device Suspension for the `soc` and `pln` Drivers” on page 18.
 - If you have suspend-unsafe devices, specify those devices in the unsafe driver list, which blocks a quiesce from starting.
For details, see “To Specify an Unsafe Driver List” on page 18.
 - If you have tape devices that are not supported by Sun Microsystems, make those devices detach-safe.
For details, see “To Make an Unsupported Tape Device Detach-Safe” on page 19.
3. If you want to use multipathing, configure your domain for multipathing and run the appropriate multipathing software on the domain.
4. Reboot the domain to process the configuration changes.

Note – You must reboot the domain after any changes to the DR configuration. If you want to minimize the number of reboots, you may want to perform various configuration tasks then reboot the domain.

5. After the reboot completes successfully, review the `/var/adm/messages` file for messages that verify the DR configuration changes.

For example, if you enabled the kernel cage, the following message is generated:

```
NOTICE: DR Kernel Cage is Enabled
```

▼ To Enable the Kernel Cage

A caged kernel confines the nonpageable memory to a minimal (most often one) number of systems boards. By default the kernel cage is disabled, preventing DR detach operations. If you plan to perform DR detach operations, you must enable the kernel cage by using the `system(4)` variable `kernel_cage_enable`, as explained in the following procedure.

Be aware that DR attach or `addboard` operations are enabled by default, regardless of the setting of the `kernel_cage_enable` variable.

Note – Before the release of version 7 of the Solaris software, the `dr-max-mem` variable was used to enable DR. The `dr-max-mem` variable is not used to enable DR in version 7 and subsequent versions of the Solaris software.

1. **Using a text editor, edit the `/etc/system` file for the domain so that `kernel_cage_enable` equals 1.**

```
set kernel_cage_enable=1
```

2. **After all DR configuration tasks are complete, be sure to reboot the domain so that this configuration can take effect.**
3. **Verify this configuration change in the `/var/adm/messages` file.**

The following example is a section from a `messages` file, which indicates that the kernel cage was enabled:

```
NOTICE: DR Kernel Cage is Enabled
```

▼ To Set Permanent Driver Parameters for Network Drivers

If you use the `ndd(1M)` command to set the driver configuration parameters for network devices, those parameters may not persist after a DR operation.

- If you want to set the driver configuration parameters permanently, set the parameters in the `/etc/system` file or the `driver.conf` file for a specific driver.

▼ To Enable Device Suspension for the `soc` and `pln` Drivers

If your system boards contain `soc` and `pln` devices, perform the following steps to make those drivers suspend-safe.

1. **Using a text editor, edit the `/etc/system` file so that the `pln_enable_detach_suspend` and `soc_enable_detach_suspend` variables are set to 1, as in the following example:**

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

2. **After all DR configuration tasks are complete, reboot the domain so that this configuration can take effect.**

▼ To Specify an Unsafe Driver List

You can provide the Solaris operating environment with information regarding any suspend-unsafe devices in the system by specifying an unsafe-driver list in the `ngdr.conf` file.

DR reads this list when it prepares to suspend the operating environment so that a board containing nonpageable memory can be detached. If DR finds an active driver in the unsafe driver list, it aborts the operation and returns an error message. The message identifies the active, unsafe driver. You must manually suspend the device so that the DR operation can be performed.

1. **Using a text editor, edit the `/platform/SUNW,Ultra-Enterprise-10000/kernel/drv/ngdr.conf` file and specify suspend-unsafe device drivers as shown below:**

```
unsupported-io-drivers="driver1","driver2","driver3";
```

where `driverx` represents each suspend-unsafe device driver.

2. **After all DR configuration tasks are complete, reboot the domain so that this configuration can take effect.**

▼ To Make an Unsupported Tape Device Detach-Safe

For the Solaris 9 operating environment, tape devices that are natively supported by Sun Microsystems are suspend-safe and detach-safe. For details, refer to the `st(7D)` man page for a list of natively-supported drives. If a system board to be detached contains a natively-supported tape device, you can safely detach the board without suspending the device.

However, if you want to use a tape device that is not natively supported by Sun Microsystems, you can use it, but you must make it detach-safe by performing the following steps.

1. Edit the `/kernel/drv/st.conf` file with the `ST_UNLOADABLE (0x0400)` flag set in the entry. Refer to the `st(7D)` man page for more information.
2. After all DR configuration tasks are complete, be sure to reboot the domain so that this configuration can take effect.

Preparing for DR Detach Operations

You must prepare a board for DR detach operations by following the steps described below. Although the following list of tasks implies a sequence of order, strict adherence to the order is not necessary. These steps apply to boards containing I/O or non-network devices. -

1. Unmount file systems.

For example, if you are using Solstice DiskSuite metadevices, you must unmount the file system for metadevices that have a board-resident partition (for example, `umount /partit`).

If you have suspend-unsafe devices that manage file systems, unmount those file systems before a detach operation. If have to manually suspend unsafe devices that manage file systems, lock those file systems using the `lockfs(1M)` command before manually suspending the unsafe devices.



Caution – Unmounting shared file systems by using the `share(1M)` utility may affect NFS client systems.

2. Remove disk partitions from the swap configuration by using `swap(1M)`.

3. If you want to detach a board that hosts Sun StorEdge A3000 controllers, make those controllers idle or take them offline manually using the `rm6` or `rdacutil` programs.

The Sun StorEdge A3000 (formerly known as the RSM Array 2000) has dual controller paths with automatic load balancing and automatic failover functionality.

4. Close all non-network devices by doing the following:
 - Close all instances of a device by killing any processes that directly open the device or raw partition, or by directing the process to close an open device on the board.
 - Run `modunload(1M)` to unload each detach-unsafe or loaded device driver.

Note – In situations where you cannot unload a device that has an unsafe driver, you can blacklist the board that contains the unsafe device and then reboot the domain. You can remove the board later. For details on blacklisting, refer to the `blacklist(1M)` man page.

5. Processes bound to the processors of a board prevent that board from being detached. You can use `pbind(1M)` to rebind them to other processors.

Configuration Changes During DR Operations

This section describes

- How you can control forcible conditions that affect system quiescence during a DR detach operation in progress
- Various configuration changes performed by DR during DR detach operations

Controlling Forcible Conditions that Affect System Quiescence

If the Solaris operating environment cannot quiesce during a DR detach operation involving a board with nonpageable memory, it displays the reason why it cannot quiesce. For example, a suspend-unsafe device is open that cannot be quiesced by the operating environment.

A failure to quiesce due to open suspend-unsafe devices is known as a *forcible condition*. You have the option to retry the operation, or you can try to force the quiescence. The conditions that cause processes not to suspend are generally temporary in nature. You can retry the operation until the quiescence succeeds.

When you try to force the quiescence, you give the operating environment permission to continue with the quiescence even if forcible conditions are still present. Doing this forces the operating environment to permit the detach. Note that, although a detach can be forced to proceed when there are open suspend-unsafe devices in the system, it is not possible to force a detach when a detach-unsafe device resides on the board and its driver is loaded

Note – Real-time processes do not prevent the operating environment from quiescence.

The most straightforward way to quiesce a domain is to close any suspend-unsafe devices. For each network driver you must execute the `ifconfig(1M)` command with its `down` parameter, then again with its `unplumb` parameter (refer to the `ifconfig(1M)` man page for more information).

Note – It should be possible to unplumb all network drivers. However, this action is rarely tested in normal environments and may result in driver error conditions. If you use DR, Sun Microsystems suggests that you test these driver functions during the qualification and installation phases of any suspend-unsafe device.

If a suspend-unsafe device is open and cannot be closed, you can manually suspend the device, and then force the operating environment to quiesce. After the operating environment resumes, you can manually resume the device as explained below.

Note – If you cannot make a device suspend its access to the domain centerplane, do not force the operating environment to quiesce. Doing so could cause a domain to crash or hang. Instead, postpone the DR operation until the suspend-unsafe device is no longer open.

▼ To Manually Suspend a Suspend-Unsafe Device

1. **Remove the usage of the device by performing one or more of the following:**
 - a. **Close the device by killing the processes using the device.**
 - b. **Ask users not to utilize the device.**

c. **Disconnect the cables to the device.**

For example, if a device that allows asynchronous unsolicited input is open, you can disconnect its cables prior to quiescing the operating environment, preventing traffic from arriving at the device and the device from accessing the domain centerplane. You can reconnect the cables after the operating environment resumes.

d. **Unload the device driver by using the `modunload(1M)` command.**

2. **Perform the DR operation again.**

3. **Do the following:**

a. **Reload the device by using the `modload(1M)` command.**

b. **Reconnect the cables to the device.**

c. **Inform users that the device can be used again.**

d. **Restart any processes associated with the device.**



Caution – If you attempt a forced quiesce operation while activity is occurring on a suspend-unsafe device, the domain may hang. However, if the domain hangs, it will not affect other domains that are running on the Sun Enterprise 10000 system.

▼ To Force a System Quiescence



Caution – Exercise care when using the force option. To successfully force the operating environment to quiesce, you must first manually quiesce the controller. Procedures to do that, if any, are device-specific. The device must not transfer any data, reference memory, or generate interrupts during the operation. Be sure to test any procedures used to quiesce the controller while it is open before running them on a production system. Using the force option to quiesce the operating environment, without first successfully quiescing the controller, can result in a domain failure and subsequent reboot.

- **For Solaris 9 (DR model 3.0) operations, run the `deleteboard(1M)` or `moveboard(1M)` command with the `-f` option.**

Target Memory Constraints

When detaching a board with nonpageable memory, DR locates an alternative (target) memory board to which to copy the nonpageable memory.

If no target board is found for a copy rename operation, the `deleteboard(1M)` and `moveboard(1M)` commands display the following error messages, respectively:

```
deleteboard: unconfigure SB2: No available memory target: dr@0:SB2::memory
```

```
moveboard: unconfigure SB2: No available memory target: dr@0:SB2::memory
```

Processors

The boot processor is responsible for maintaining the `netcon` BBSRAM buffer.

Before detaching a board on which the boot processor resides, DR assigns the boot processor role to another active (online) processor.

Network Devices

A detach operation will fail if there are any network interfaces on the board that fit the following conditions. In these cases, the detach operation fails and DR displays an error message.

- The interface is the primary network interface for the domain; that is, the interface whose IP address corresponds to the network interface name contained in the file `/etc/nodename`.

Note that bringing down the primary network interface for the domain 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 on the same subnet as the SSP host for the system; that is, the subnet of the IP address that corresponds to the SSP host name found in `/etc/ssphostname`.

Bringing down this interface interrupts communication between the host and SSP. Because DR operations are initiated on the SSP, control of the detach process would be lost. Note that the `/etc/ssphostname` file contains the name of the SSP that controls the host; therefore, if you rename the SSP, you must manually update the `/etc/ssphostname` file.



Caution – Detaching network interfaces may affect NFS client systems.

Remote DR Communication

In Solaris 9 domains, the domain configuration server, `dcS(1M)`, controls DR operations.

▼ To Troubleshoot a Connection Failure During a Solaris 9 (DR Model 3.0) Operation

1. Check the domain.

`dcS(1M)` must be configured in the `/etc/inetd.conf` file of the domain. The following lines must be present in the file:

```
sun-dr stream tcp    wait root /usr/lib/dcs dcs
sun-dr stream tcp6  wait root /usr/lib/dcs dcs
```

2. If the `dcS` daemon is configured in `/etc/inetd.conf`, kill `dcS(1M)` if it is currently running, and send a HUP signal to the `inetd(1M)` daemon to cause it to re-read the `inetd.conf(4)` configuration file:

```
# kill -9 dcs_pid
# kill -HUP inetd_pid
```

where `dcS_pid` is the process ID of the `dcS(1M)` daemon and `inetd_pid` is the process ID of the `inetd(1M)` daemon.

3. Check the `/var/adm/messages` file for possible error messages from `inetd(1M)` if it is having trouble starting `dcS(1M)`.

The `dcS(1M)` daemon executable file resides in the `/usr/lib` directory.

4. Try the DR operation again, starting from the beginning.

Index

A

addboard(1M), 3

C

cfgadm(1M), 3

commands

addboard(1M), 3

cfgadm(1M), 3

deleteboard(1M), 3

moveboard(1M), 3

rcfgadm(1M), 3

showdevices(1M), 3

configuring swap space I/O controllers across boards, 5

connection, loss of, 14

D

dcs(1M), 14

deleteboard(1M), 3

detach

closing devices, 9

pageable memory and swap space during detach, 5

RSM 2000 and detach, 10

Sun StorEdge A3000 and detach, 10

swap space and detach, 5

detach and processors, 13

detach-safe tape devices, 9

devices

detach-safe, 4

detach-unsafe, 4, 11

preparing for detach, 9

suspend-safe, 4

suspend-unsafe, 4

disk swap space, and detach, 5

DR

where to execute commands, 3

DR models

2.0 model, 2

3.0 model, 2

drivers, suspend-safe, 5

F

files

st.conf (ST_UNLOADABLE flag and tape devices), 9

forcible conditions and quiesce failures, 11

L

loss of connection, 14

M

manually suspending suspend-unsafe devices, 11

moveboard(1M), 3

multipathing, 3

N

network drivers, suspend-unsafe, 11

tape devices, detach-safe, 9

tape devices, suspend-unsafe, 9

timeout, RPC, 14

P

pageable memory and swap space, during
detach, 5

processors and detach, 13

Q

quiescence, 4

failure reasons, 10

forcible conditions, 11

suspend-unsafe devices, 10

R

rcfgadm(1M), 3

RPC timeout, 14

RSM 2000 and detach, 10

S

showdevices(1M), 3

ST_UNLOADABLE flag and tape devices, 9

Sun StorEdge

A3000 and detach, 10

suspend failures and forcible conditions, 11

suspend, reasons it may fail, 10

suspending OS and suspend-unsafe devices, 10

suspend-safe drivers listing, 5

suspend-unsafe device and quiescing OS, 10

suspend-unsafe devices, manually suspending, 11

suspend-unsafe tape devices, 9

swap space, configuring for detach, 5

swap space, configuring I/O controllers across
boards, 5

T

tape devices and ST_UNLOADABLE flag, 9