



Sun Enterprise™ 10000 DR Configuration Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900 U.S.A.
650-960-1300

Part No. 816-2095-10
October 2001, Revision A

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

Copyright 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 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 2001 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, CA 94303-4900 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.



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 vii

Before You Read This Book vii

How This Book Is Organized vii

Using UNIX Commands viii

Typographic Conventions ix

Shell Prompts ix

Related Documentation x

Accessing Sun Documentation Online x

Ordering Sun Documentation xi

Sun Welcomes Your Comments xi

1. Configuring DR 1

DR Models 2

DR Model 2.0 2

DR Model 3.0 3

Getting Started 4

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	8
▼ To Enable Device Suspension for the <code>soc</code> and <code>pln</code> Drivers	8
▼ To Specify an Unsafe Driver List	8
▼ To Make an Unsupported Tape Device Detach-Safe	9
▼ To Select a DR Model	10
Preparing for DR Detach Operations	10
Configuration Changes During DR Operations	12
Controlling Forcible Conditions that Affect System Quiescence	12
▼ To Manually Suspend a Suspend-Unsafe Device	13
▼ To Force a System Quiescence	14
Target Memory Constraints	14
Processors	15
Network Devices	15
Remote DR Communication	16
▼ To Troubleshoot an RPC Time-out or Connection Failure During a DR Model 2.0 Operation	17
▼ To Troubleshoot a Connection Failure During a DR Model 3.0 Operation	17
DR Model 2.0 Configuration Details	18
Reconfiguration After a DR Model 2.0 Operation	18
When to Reconfigure the Domain	18
Disk Devices	19
DR Model 2.0 and AP Interaction	20

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 x.

Before You Read This Book

This guide is intended for the dynamic reconfiguration 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 all of the books in the Solaris System Administration collection in AnswerBook2™ format provided with your server 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 in AnswerBook2 format.

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 information on basic UNIX commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following sources for this information:

- AnswerBook2 online documentation for the Solaris software environment, particularly those dealing with Solaris system administration
- Other software documentation that you received with your system

Typographic Conventions

Typeface or Symbol	Meaning	Examples
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. % You have mail.
AaBbCc123	What you type, when contrasted with on-screen computer output	% su Password:
<i>AaBbCc123</i>	Book titles, new words or terms, words to be emphasized. 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>	806-7616
	<i>Sun Enterprise 10000 SSP 3.5 User Guide</i>	806-7613
	<i>Sun Enterprise Server Alternate Pathing 2.3.1 User Guide</i>	806-4150
	<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>	806-4121
Reference	<i>Sun Enterprise 10000 Dynamic Reconfiguration Reference Manual</i>	808-7617
	<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>Sun Enterprise Server Alternate Pathing 2.3.1 Reference Manual</i>	806-4151
	<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>

Ordering Sun Documentation

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

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

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

Sun Welcomes Your Comments

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

docfeedback@sun.com

Please include the part number (816-2095-10) 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
- Information specific to DR model 2.0
 - How to reconfigure DR model 2.0 after a DR operation on domains running the Solaris 7, 2.6, or 2.5.1 operating environments
 - How DR model 2.0 interacts with the Sun Enterprise Server Alternate Pathing feature

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 Hostview, `dr` shell, or the ADR commands. For instructions on detaching boards from DR model 2.0 or 3.0 domains, refer to the *Sun Enterprise 10000 Dynamic Reconfiguration User Guide*.

DR Models

With the release of the Solaris 8 10/01 operating environment, two different models for dynamic reconfiguration are supported on Sun Enterprise 10000 domains. Both these models, referred to as DR model 2.0 and DR model 3.0, enable you to logically attach and detach system boards to and from the Solaris operating environment with minimum domain downtime.

DR is used in conjunction with hot swap, which is the process of physically removing or inserting a system board. Use DR to do the following:

- Add a system board.
- Reinstall a repaired system board.
- Move a system board from one domain to another.
- Delete a system board.
- Modify the domain configuration on a Sun Enterprise 10000 system.

Note that only one DR model can run on a domain at a time. The following sections describe the differences between DR model 2.0 and 3.0.

DR Model 2.0

DR model 2.0 is the default DR model on Sun Enterprise 10000 domains. This model uses the `dr_daemon(1M)` to control DR operations on domains. You can use the following features on the System Service Processor (SSP) to perform DR operations:

- Hostview
- `dr(1M)` shell
- ADR commands `addboard(1M)`, `deleteboard(1M)`, `moveboard(1M)`, and `showusage(1M)`

For details on using Hostview, refer to the *Sun Enterprise 10000 SSP 3.5 User Guide* and the *Sun Enterprise 10000 Dynamic Reconfiguration User Guide*. For information on using `dr` shell and the ADR commands, refer to the *Sun Enterprise 10000 Dynamic Reconfiguration User Guide*.

Note that DR model 2.0 is the default DR model in all releases of the Solaris 8, 7, 2.6, and 2.5.1 operating environments.

If you intend to use or are currently using alternate pathing software on your domains, note that Sun Enterprise Server Alternate Pathing software must be used with DR model 2.0 domains. If you are running the Solaris 8 operating environment on your domain, you must use AP 2.3.1.

For information on how DR model 2.0 and alternate pathing work together, see “DR Model 2.0 and AP Interaction” on page 20. For details on alternate pathing, refer to the *Sun Enterprise Server Alternate Pathing User Guide*.

DR Model 3.0

DR model 3.0 is new in this release of the Solaris 8 10/01 operating environment. DR model 3.0 uses the domain configuration server, `dcs(1M)`, to control DR operations on Sun Enterprise 10000 domains. You use the ADR commands `addboard(1M)`, `moveboard(1M)`, and `deleteboard(1M)` to perform DR operations and the `showdevices(1M)` and `rcfgadm(1M)` commands to obtain device and board status information. You run these commands on the SSP. For details on performing DR model 3.0 operations, refer to the *Sun Enterprise 10000 Dynamic Reconfiguration User Guide*.

DR model 3.0 domains also interface with the Reconfiguration Coordination Manager (RCM), which enables you to coordinate DR operations with other applications, such as database, clustering, and volume management software running on a domain. For details on RCM, refer to the *Solaris 8 System Administration Supplement* in the Solaris 8 10/01 Update Collection.

Note that DR model 3.0 is available only in the Solaris 8 10/01 release of the Solaris operating environment. If you are running DR model 3.0, you must also run SSP 3.5 software on the SSP.

If you intend to use or are currently using multipathing software on your DR model 3.0 domains, use IPMP (the IP multipathing software provided with the Solaris operating environment) and the Sun StorEdge™ Traffic Manager (also called MPxIO).

For details on multipathing, refer to the *IP Network Multipathing Administration Guide* in the Solaris 8 10/01 Update Collection and the *MPxIO Installation and Configuration Guide*, available through the Sun Download Center (<http://www.sun.com/download>). For information on obtaining MPxIO software and documentation, refer to the *SSP 3.5 Installation Guide and Release Notes*.

Note – MPxIO may not support automatic path switching for all devices. Refer to the *MPxIO Installation and Configuration Guide* for details. If you require automatic path switching for devices that are not supported, use Alternate Pathing software with DR model 2.0.

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 4.
- Verify that you have sufficient swap space for your domain.
For details, see “Allocating Sufficient Domain Swap Space” on page 5.
- Qualify any third-party device drivers, as described in “Qualifying Third-Party Device Drivers” on page 5.

Device Prerequisites

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

- *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 10.

- *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 13.

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`, `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`, `ses`).

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 either DR model 2.0 or 3.0 domains. 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 7.
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 8.
 - 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 8.
 - 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 8.
 - 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 9.
3. Determine which DR model you want to run on the domain and if necessary switch the DR model, as discussed in “To Select a DR Model” on page 10.
4. If you want to use multipathing, configure your domain for multipathing and run the appropriate multipathing software on the domain.
For information on the multipathing software compatible with each DR model, see “DR Model 2.0” on page 2 and “DR Model 3.0” on page 3.
5. 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.

6. 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 and switched the DR model from 2.0 to 3.0, the following messages are generated:

```
NOTICE: DR Kernel Cage is Enabled
.
.
.
NOTICE: Next Generation DR Model (3.0) 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 the Solaris 7 operating environment, the `dr-max-mem` variable was used to enable DR. This variable is no longer used in the Solaris 7 and Solaris 8 operating environments.

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 `dr.conf` file (for DR model 2.0 domains) and `ngdr.conf` file (for DR model 3.0 domains).

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 following files and specify suspend-unsafe device drivers as shown below:

- /platform/SUNW,Ultra-Enterprise-10000/kernel/drv/dr.conf
- /platform/SUNW,Ultra-Enterprise-10000/kernel/drv/ngdr.conf

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

where *driver* is the name of the suspend-unsafe device driver.

Note – It is suggested that you list all unsafe drivers in both the `dr.conf` and `ngdr.conf` files, whether or not you intend to switch between DR models. If you decide later to switch DR models, both configuration files will include references to all unsafe drivers.

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 8 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.**

▼ To Select a DR Model

Switching a domain from DR model 2.0 to model 3.0 or from model 3.0 to model 2.0 involves modifying the `ngdr.conf` file, then rebooting the domain.

- **Edit the `ngdr.conf` file as described in the *Sun Enterprise 10000 Dynamic Reconfiguration User Guide*.**

Note – If you want to switch from DR model 2.0 to model 3.0 on a domain, be aware that DR model 3.0 domains must run the Solaris 8 10/01 operating environment. The SSP must also be running SSP 3.5 software.

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. Be aware that some steps apply to DR model 2.0 only, so if you are running DR model 3.0, skip those steps.

1. For DR model 2.0 domains, note the following if you are detaching a board with non-network devices:
 - If you are using the redundancy features of Alternate Pathing or Solstice DiskSuite™ mirroring to access a non-network device connected to the board, reconfigure these subsystems so that the device or network is accessible using controllers on other system boards.

Alternate Pathing software automatically switches the disk devices to an alternate interface if one is available.
 - Remove Alternate Pathing or Solstice DiskSuite databases from board-resident partitions. The location of Alternate Pathing or Solstice DiskSuite databases is explicitly chosen by the user and can be changed.
 - Remove any private regions used by Sun Enterprise Volume Manager™ or Veritas Volume Manager.

Volume manager by default uses a private region on each device that it controls, so such devices must be removed from volume manager control before they can be detached.
2. 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.

3. Remove disk partitions from the swap configuration by using `swap(1M)`.
4. 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.

5. 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.

If you are using DR model 2.0, you can use the Hostview device display or the `drshow(1M)` I/O listing to view the open count field, which indicates how many processes have opened particular devices. Use the `fuser(1M)` command on the domain to see which processes have these devices open.
 - 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.

6. If the Solaris operating environment must be suspended, kill all the real-time processes that are open or unbind (kill) any process bound to an on-board processor.

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 its reasons, which may include the following:

- Real-time processes are running in the domain.
- A device that cannot be quiesced by the operating environment (that is, a suspend-unsafe device) is open.

A failure to quiesce due to real-time processes or 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 (real-time processes are running or a suspend-unsafe device is open) 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

If a real-time process is running, determine if suspending the process would produce an adverse effect on the functions performed by the process. If not, you can force the operating environment to quiesce.

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.

1. For DR model 2.0 operations, do one of the following:

Click the `Force` button within Hostview as described in the *Sun Enterprise 10000 Dynamic Reconfiguration Guide*.

- Enter the `complete_detach(1M)` command with the `force` option within the `dr(1M)` shell application
- Run the `deleteboard(1M)` or `moveboard(1M)` command with the `-f` option.

2. For 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.

In DR model 2.0, if no target board is found, the detach operation is refused, and DR displays the following warning message on the system console:

```
WARNING: sfdr: sfdr_pre_release_mem: no available target for mem-  
unit (board.0)
```

In DR model 3.0, if no target board is found for a copy rename operation, the `deleteboard(1M)` and `moveboard(1M)` commands display the following error messages:

```
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

For DR model 2.0 detach operations, DR automatically terminates usage of all network interfaces on the board that is being detached. When you complete the detach operation, the `dr_daemon(1M)` identifies all configured interfaces on the board being detached and issues the following `ifconfig(1M)` commands on each such interface.

```
ifconfig interface down  
ifconfig interface unplumb
```

Additionally, if FDDI interfaces are detached, DR kills the FDDI network monitoring daemon before you perform the detach operation. DR then restarts it after the detach is complete. Note that the `/usr/sbin/nf_snmd` daemon for `nf` devices is neither started nor stopped when a board that contains a FDDI interface is attached.

DR does not execute these commands on a board that contains a network interface that fits any of 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.

- The interface is the active alternate for an Alternate Pathing (AP) metadvice when the AP metadvice is plumbed. Interfaces used by AP should not be the active path when the board is being detached.

AP 2.1 performs the switch automatically; however, you can manually switch the active path to an interface that is not on the board being detached. If no such path exists, manually execute the `ifconfig down` and `ifconfig unplumb` commands on the AP interface. To manually switch an active path, use the `apconfig(1M)` command.



Caution – Detaching network interfaces may affect NFS client systems.

Remote DR Communication

In DR model 2.0 domains, the `dr_daemon(1M)` communicates with Hostview and the `dr(1M)` shell application (both of which run on the SSP) by way of Remote Procedure Calls (RPCs). In DR model 3.0 domains, the domain configuration server, `dcs(1M)`, controls DR operations.

If you experience connection failures during a DR operation, perform the procedure appropriate to the DR model running on your domain.

▼ To Troubleshoot an RPC Time-out or Connection Failure During a DR Model 2.0 Operation

1. Check the domain.

The daemon must be configured in the `/etc/inetd.conf` file of the domain. The following line (which appears on a single line) must be present in the file:

```
300326/4 tli rpc/tcp wait root \  
/platform/SUNW,Ultra-Enterprise-10000/lib/dr_daemon/ dr_daemon
```

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

```
# kill dr_daemon_pid  
# kill -HUP inetd_pid
```

where `dr_daemon_pid` is the process ID of the `dr(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 the `dr_daemon(1M)`.

The DR daemon executable file resides in the `/platform/SUNW,Ultra-Enterprise-10000/lib` directory.

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

▼ To Troubleshoot a Connection Failure During a 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 `dc`s daemon is configured in `/etc/inetd.conf`, kill `dc`s(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 `dc`s(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 `dc`s(1M).
The `dc`s(1M) daemon executable file resides in the `/usr/lib` directory.
4. Try the DR operation again, starting from the beginning.

DR Model 2.0 Configuration Details

This section provides configuration information specific to DR model 2.0.

Reconfiguration After a DR Model 2.0 Operation

You must reconfigure your DR model 2.0 domain after you have attached or detached a system board. DR model 2.0 enables you to reconfigure the domain after a DR operation.

Note – If your DR model 2.0 domain is running the Solaris 8 operating environment, you no longer need to perform the manual reconfiguration procedures described in this section. A new DDI subsystem, `devfsadm`, completes all the reconfiguration tasks. For details, see the `devfsadm`(1M) man page.

When to Reconfigure the Domain

You must reconfigure the domain if any of the following conditions occur:

- Board addition – When you add a board to a domain, you must execute the reconfiguration sequence to configure the I/O devices that are associated with the board.

- Board deletion – If you remove a board that is not to be replaced, you may, but do not have to, execute the reconfiguration sequence to clean up the `/dev` links.
- Board replacement – If you remove a board then reinsert it in a different slot or if you replace a board with another board that has different I/O devices, you must execute the reconfiguration sequence to configure the I/O devices that are associated with the board. However, if you replace a board with another board that hosts the *same* set of I/O devices, inserting the replacement into the *same* slot, you do not need to execute the reconfiguration sequence. But, be sure to insert a replacement board into the same slot that was vacated to retain the original mapping of `/dev` links to physical names.

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

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

When you execute the reconfiguration sequence after you attach a board, device path names not previously seen by the domain are written to the `/etc/path_to_inst` file. The same path names are also added to the `/devices` hierarchy, and links to them are created in the `/dev` directory.

Disk Devices

Disk controllers are numbered consecutively as the `disks(1M)` program encounters them. All disk partitions are assigned `/dev` names according to the disk controller number that `disks(1M)` assigns. For example, all disk partitions that are accessible using disk controller 1 are named `/dev/dsk/cXtYdZsW`

where:

`x` is the disk controller number,

`Y`, in most cases, corresponds to the disk target number,

`Z` corresponds to the logical unit number, and

`W` corresponds to the partition number.

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

Note – The disk controller number is part of the `/dev` link name used to access the disk. If that number changes during the reconfiguration sequence, the `/dev` link name also changes. This change may affect file system tables and software, such as Solstice DiskSuite, which use the `/dev` link names. Update `/etc/vfstab` files and execute other administrative actions to change the `/dev` link names.

DR Model 2.0 and AP Interaction

A DR detach operation works with Alternate Pathing (AP) or Solstice DiskSuite mirroring when you detach a board that hosts I/O controllers that are attached to vital system resources. If, for example, the root (`/`) or `/usr` partition is on a disk attached to a controller on the board, the board cannot be detached unless there is a hardware alternate path to the disk, and AP has been configured to take advantage of it, or the disk is mirrored. The alternate path or the mirrors must be hosted by other boards in the domain. The same applies to network controllers. The board that hosts the Ethernet controller that connects the SSP to the Sun Enterprise 10000 platform cannot be detached unless an alternate path exists to an Ethernet controller on another board for this network connection.

DR notifies the AP subsystem when system boards are attached, detached, or placed in the drain state. In addition, DR queries AP about which controllers are in the AP database and their status (active or inactive). This communication occurs between the `dr_daemon(1M)` and `ap_daemon(1M)`. If the `ap_daemon(1M)` is not present, an error message is placed in the syslog messages buffer of the domain and DR operations continue without error.

To disable the interaction between DR and the `ap_daemon`, use the `-a` option when you invoke `dr_daemon(1M)`. Refer to the `dr_daemon(1M)` command in the *Sun Enterprise 10000 Dynamic Reconfiguration Reference Manual*.

If you are using AP version 2.1, the operating environment automatically switches off the active disk controllers on outgoing boards during the complete-detach phase of DR. If you are using AP version 2.0, you need to manually switch off the active disk controllers before you start the complete-detach phase. If your domain is running the Solaris 8 operating environment, you must use AP version 2.3.1.

For more information about DR and AP interaction, see the *Sun Enterprise Servers Alternate Pathing 2.3.1 User Guide*. For more information about AP and Solstice DiskSuite, refer to the *RAS Companion*.

Index

A

- alternate pathing (AP)
 - interaction with DR, 20
 - Solstice DiskSuite, 20
 - vital partitions during detach, 20
- attach
 - reconfiguration sequence after attach, 19
- automatically switching off active controllers during detach, 20

B

- board added, reconfigure after, 18
- board deleted, optionally reconfigure after, 19
- board replaced, reconfigure after, 19

C

- communication timeouts affecting Hostview and dr(1M), 16
- configuring swap space I/O controllers across boards, 5
- connection, loss of, 16
- controllers (disk), number of, 19

D

- dcs(1M), 3, 16
- detach

- closing devices, 10
- I/O controllers on board being detached, 20
- network between SSP and UE10000, and detach, 20
- network controllers and detach, 20
- pageable memory and swap space during detach, 5
- RSM 2000 and detach, 11
- Sun StorEdge A3000 and detach, 11
- swap space and detach, 5
- switching of active controllers during detach, 20
- detach and network devices, 15
- detach and processors, 15
- detach-safe tape devices, 9
- dev, reconfiguring /dev links after DR operation, 18
- devices
 - detach-safe, 4
 - detach-unsafe, 4, 12
 - network, 15
 - preparing for detach, 10
 - suspend-safe, 4
 - suspend-unsafe, 4
- disabling AP / DR interaction, 20
- disk controller numbering, 19
- disk devices, reconfiguring after DR operation, 19
- disk swap space, and detach, 5
- DR / AP interaction, disabling, 20
- DR models, 2
 - 2.0 model, 2, 14, 15, 18
 - 3.0 model, 3, 15
- dr_daemon(1M), 2, 16

drivers, suspend-safe, 5

E

Ethernet between SSP and UE10000, and detach, 20

F

files

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

forcible conditions and quiesce failures, 12

I

I/O controllers on board being detached, 20

I/O devices, configuring for detach, 20

I/O devices, reconfiguring after DR operation, 18

IPMP (IP multipathing), 3

L

loss of connection, 16

M

manually suspending suspend-unsafe devices, 13

MPxIO, 3

N

network between SSP and UE10000, and detach, 20

network controllers and detach, 20

network drivers, suspend-unsafe, 12

numbering of disk controllers, 19

P

pageable memory and swap space, during detach, 5

processors and detach, 15

Q

quiescence, 4

- failure reasons, 12

- forcible conditions, 12

- forcing, 14

- real-time processes, 12

- suspend-unsafe devices, 12

R

real-time processes and quiescing OS, 12

reconfiguration

- disk devices, 19

- domains, 18

root partition and I/O controllers during detach, 20

RPC timeout, 16

RSM 2000 and detach, 11

S

Solstice DiskSuite and mirroring, 20

ST_UNLOADABLE flag and tape devices, 9

Sun StorEdge

- A3000 and detach, 11

- Traffic Manager, 3

suspend failures and forcible conditions, 12

suspend, reasons it may fail, 12

suspending OS and real-time processes, 12

suspending OS and suspend-unsafe devices, 12

suspend-safe drivers listing, 5

suspend-unsafe device and quiescing OS, 12

suspend-unsafe devices, manually suspending, 13

suspend-unsafe tape devices, 9

swap space, configuring for detach, 5

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

switching off active controllers during detach, 20

system board added, reconfigure after, 18

system board deleted, optionally reconfigure after, 19

system board replaced, reconfigure after, 19

T

tape devices and ST_UNLOADABLE flag, 9

tape devices, detach-safe, 9

tape devices, suspend-unsafe, 9

timeout, RPC, 16

timeouts affecting Hostview and dr(1M), 16

U

usr partition and I/O controllers during detach, 20

